

Closed-loop control of an experimental mixing layer using machine learning control

Vladimir Parezanović¹ †, Thomas Duriez^{1,2}, Laurent Cordier¹,
Bernd R. Noack¹, Joël Delville¹, Jean-Paul Bonnet¹,
Marc Segond³, Markus Abel^{3,4,5} and Steven L. Brunton⁶

¹Institut PPRIME, CNRS - Université de Poitiers - ENSMA, UPR 3346, Département Fluides, Thermique, Combustion, CEAT, 43, rue de l'Aérodrome, F-86036 Poitiers Cedex, France

²Laboratorio de Fluidodinamica - Facultad de Ingenieria CONICET - Universidad de Buenos Aires, Paseo Colon 850, Ciudad Autonoma de Buenos Aires, Argentina

³Ambrosys GmbH, Albert-Einstein-Str. 1-5, D-14469 Potsdam, Germany

⁴LEMTA, 2 avenue de la Forêt de Haye, F-54518 Vandoeuvre-lès-Nancy Cedex, France

⁵University of Potsdam, Karl-Liebknecht-Str. 24/25, D-14476 Potsdam, Germany

⁶University of Washington, Applied Math Department, Seattle, WA 98195, USA

(Received ?; revised ?; accepted ?. - To be entered by editorial office)

A novel framework for closed-loop control of turbulent flows is tested in an experimental mixing layer flow. This framework, called Machine Learning Control (MLC), provides a model-free method of searching for the best function, to be used as a control law in closed-loop flow control. MLC is based on genetic programming, a function optimization method of machine learning. In this article, MLC is benchmarked against classical open-loop actuation of the mixing layer. Results show that this method is capable of producing sensor-based control laws which can rival or surpass the best open-loop forcing, and be robust to changing flow conditions. Additionally, MLC can detect non-linear mechanisms present in the controlled plant, and exploit them to find a better type of actuation than the best periodic forcing.

1. Introduction

Closed-loop turbulence control is fast gaining in significance as a research topic in fluid mechanics. Engineering benefits of such control applied to a realistic flow can be enormous. Some examples include: reduction of CO_2 emissions and fuel consumption optimization for large transport vehicles through drag reduction, optimization of green energy harnessing from wind and water turbines by continuous control of lift and drag properties of turbine blades, countless opportunities for medical applications, etc. Among all the control strategies, feedback control has been extremely successful in fluid dynamics. For many applications, such as the stabilization of unstable steady-states, linear model-based feedback control suppresses instabilities and keeps the flow approximately linear so that models and controllers remain effective. Some examples include the transitional channel flow (Ilak & Rowley 2008; Ilak 2009), flows of backward-facing step (Hervé *et al.* 2012), the flat plate boundary layer (Bagheri *et al.* 2009; Semeraro *et al.* 2013), and cavity-flow oscillations (Rowley 2002; Illingworth *et al.* 2010), although there are many others. Balanced reduced-order models have been especially effective because they capture the most relevant flow states that are both controllable and observable, based on input-output data from simulations or experiments (Rowley 2005; Ma *et al.* 2011).

† Email address for correspondence: vladimir.parezanovic@univ-poitiers.fr

For fully turbulent flows, however, most controllers rely on either open-loop or model-free adaptive control strategies. This is largely due to the fact that turbulent flows are characterized by strongly nonlinear dynamics that evolve on a high-dimensional attractor. For these flows, linear models are incapable of capturing many significant mechanisms, including broadband frequency cross-talk. Some successful studies of wake stabilization with high-frequency actuation (Glezer *et al.* 2005; Thiria *et al.* 2006; Luchtenburg *et al.* 2009) and low-frequency forcing (Pastoor *et al.* 2008) show that frequency cross-talk can be a crucial potential target for actuation. Moreover, the goal of control in a turbulent system may not be the stabilization of a laminar solution, but rather the enhancement of turbulent energy. In the appendix, we investigate the performance of balanced linear models of the mixing layer experiment and illustrate the severe limitations.

The adaptive control approaches are mostly based on a slow variation of amplitude and frequency of periodic actuation until continuously monitored performance reaches a maximum (King 2010). Examples of such approaches include resonance frequency adaptation using extremum seeking (Beaudoin *et al.* 2006), and amplitude selection using slope seeking (King *et al.* 2006). In these cases, the response of non-linear mechanisms to actuation is taken into account in performance evaluation, but the periodic forcing is, by nature, incapable of in-time targeting of specific flow events, in order to affect these non-linearities in-time directly.

A few successful examples of an in-time control are the reduction of skin friction in wall turbulence (Choi *et al.* 1994) and phasor control for turbulence with a dominant oscillatory structure (Samimy *et al.* 2007). In the former, a simple opposition control in the viscous sublayer is already effective; while in the latter, the control design requires a robust phase detection from the sensors and effective gain scheduling for the actuators. In general, a robust control-oriented reduced-order model, which could at least resolve the turbulent coherent structures and transients between them, would be required for an experimental application of a model-based control design.

Aerodynamics turns its eye once more to bird and insect flight for inspiration. A bird is a flying object whose surface is completely covered with sensors, and at the same time fully elastic so that each small part acts as an independent actuator. Seeing a bird in flight, one might think this action is effortless, but humans must think in terms of complex transfer functions between millions of sensors (skin receptors) and thousands of actuators (feathers). Numerical exploration of bio-inspired actuation already yields intriguing results (Bagheri *et al.* 2012). Development of modern actuation systems, such as micro-valves, piezoelectric, synthetic jets, plasmas, etc. (Cattafesta & Sheplak 2011; Corke *et al.* 2010; Glezer & Amitay 2002), provides progressively increasing options for application of control in experimental studies. Concurrent increase in capabilities of real-time systems with large input/output data rates and computational power bring us ever closer to having an affordable, yet powerful and compact multiple-input-multiple-output (MIMO) control system. Could we also take inspiration from the living organisms in learning how to create control laws for such complex sensor/actuator systems, in flow conditions where linear models and adaptive control do not suffice?

Recently, we have proposed Machine Learning Control (MLC) as a framework for a model-free closed-loop control of turbulent flows (Duriez *et al.* 2014; Parezanovic *et al.* 2014; Gautier *et al.* 2014). MLC is based on genetic programming (Koza 1992; Koza *et al.* 1999), a biologically-inspired machine learning function optimization method, used to design controllers in robotics (Wahde 2008; Lewis *et al.* 1992; Nordin & Banzhaf 1997). The use of machine learning for control (Fleming & Purshouse 2002) also includes genetic algorithms which can only be used to optimize control parameters (de la Fraga & Tlelo-Cuautle 2014) and artificial neural networks (Noriega & Wang 1998).

The current article will present the application of MLC in a turbulent mixing layer. Open-loop forcing of a shear layer has been proven effective at changing its stream-wise evolution (Oster & Wygnanski 1982; Fiedler 1998; Wiltse & Glezer 1993). In this experiment, actuation is applied at the trailing edge of the splitter plate with a goal of changing the local properties of the flow, such as the fluctuating energy content or thickness of the mixing layer. The sensors placed at the location of interest, downstream of the trailing edge, are subject to a convective time delay. MLC is employed in finding the best sensor-based, feedback control law, which can rival or exceed the efficiency of open-loop forcing, while ensuring robustness to variations in the mean flow conditions.

The main features of the experimental installation and the sensor/actuator system are described in § 2. Formulation of the objective functions, a detailed description of the methodology of MLC and its practical application in the experiment are described in § 3. In § 4 results of the un-actuated flow are presented, followed by a study of the mixing layer’s response to open-loop forcing, in order to have a base of comparison with closed-loop results. The performance of the best closed-loop control laws, obtained by MLC, is then presented for maximization and minimization of objective functions in two different mixing layer speed configurations. An experiment using high-gain actuation amplitude control, and the implications of its effects on the MLC process, is also discussed. In § 5, we examine different parameters and the convergence process of experimental MLC, as well as the robustness of open- and closed-loop control to varying flow conditions. The results are summarized and a discussion on MLC effectiveness and future development is provided in § 6. In Appendix A, eigensystem realization algorithm (ERA) and observer/Kalman filter identification (OKID) are performed on the measurements from the mixing layer. The results of this procedure are presented and discussed as our model-based benchmark.

2. Experimental setup

In this section, the wind-tunnel facility and the geometry of the key elements in the test section (splitter plate) are described. We present the design of the actuator system and discuss the control of the actuation amplitude and evaluation of the invested energy of actuation. Finally, sensor and acquisition systems are described.

2.1. Wind tunnel

The wind tunnel facility has been designed and built as a part of the TUCOROM† project. It features a twin-turbine design with two separate air streams, running in parallel and meeting in the test section to produce a mixing layer. The turbines are capable of propelling the air at velocities up to 12 m s^{-1} at the inlet of the test section. These velocities are however not practically obtained due to the installation of foam layers at key points along the stream paths (as shown in figure 1) whose function is to homogenize the flow and prevent the forming of large structures upstream of the test section. This was a necessary design choice as the initial stages of the wind tunnel must be compact, in order for the test section to be long enough to allow for unimpeded spatial development of the mixing layer flow. In addition, the convergent part of the wind tunnel has trip-wires installed which ensure that the boundary layers at the end of the splitter plate are turbulent for operation at higher speeds *i.e.* when a fully turbulent mixing layer is desired.

The test section is $l_x \times l_y \times l_z = 3 \text{ m} \times 1 \text{ m} \times 1 \text{ m}$, with a square cross-section. It begins

† ANR project: "TURbulence COnTrol using Reduced-Order Models"

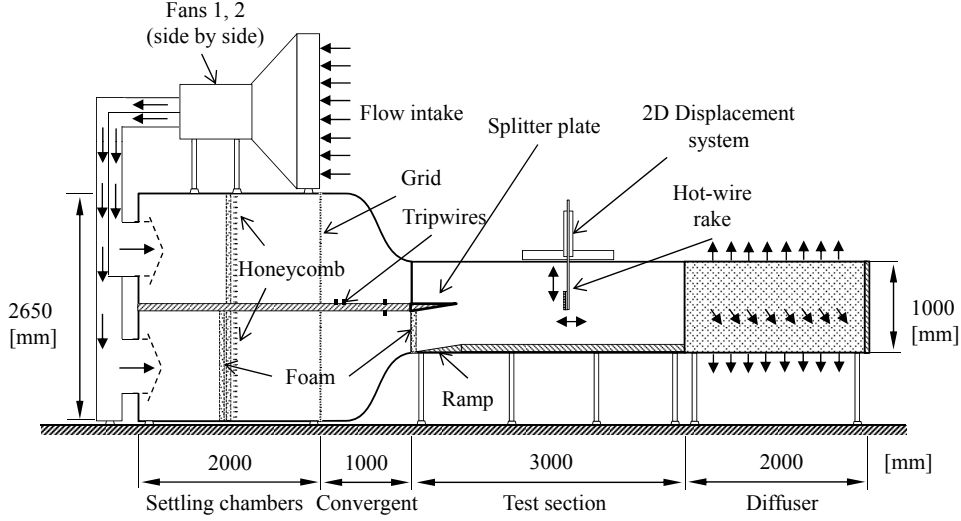


FIGURE 1. Dual stream TUCOROM wind tunnel facility schematic.

at the throat of the convergent part of the wind tunnel, where the surface separating the two air streams terminates with a splitter plate. The splitter plate maintains a horizontal upper surface (from where the upper stream originates), while its lower surface is angled at 8 degrees so as to connect with the upper surface and form a relatively thin trailing edge (see figure 2). The initial thickness of the splitter plate is 8 cm and the trailing edge is only 2 mm thick. The splitter plate spans the entire width of the test section inlet, and extends by 39 cm into the test section. The bottom of the test section has been raised by inserting a ramp of the same shape as the lower surface of the splitter plate, in order to compensate for adverse pressure gradient on the lower stream boundary layer caused by the angling of the splitter plate. This reduces the height of the test section by 8 cm. An additional foam layer is inserted at the convergent throat just upstream of the angled splitter plate in order to further stabilize the lower stream and its boundary layer. All these provisions limit the maximum operational speeds to 9.5 m s^{-1} for the upper and 1.7 m s^{-1} for the lower stream.

The test section terminates with a $2 \text{ m} \times 1 \text{ m} \times 1 \text{ m}$ foam-lined diffuser which recycles the air into the surrounding environment and prevents the structures created by the mixing layer from returning into the turbine inlets.

2.2. Actuator system

The actuator system comprises 96 micro jets at the trailing edge of the splitter plate, along its entire span. As shown in figure 2 the jets actuate in the stream-wise direction at the origin of the mixing layer. The actuators function in a binary manner; they can be commanded to open and close at a maximum rate of 500 Hz. A plenum chamber, embedded in the splitter plate, directly supplies air to the micro-jets. The amplitude of actuation is globally adjusted by control of the pressure level P_t in the plenum. For this purpose we use a PI controller which uses inputs from a differential pressure sensor, inside the plenum, and outputs command signals to an electrically controlled valve. The valve adjusts the flow rate of compressed air into the plenum chamber. The PI controller works at a rate of 100 Hz.

Due to the long tubing connection between the control valve and the plenum, the sys-




FIGURE 2. Schematic representation of the splitter plate and the actuator system layout.

tem reaction time is on the order of seconds. Therefore, the actuation amplitude can only be maintained on a long time scale average and not instantaneously. It is typically used to maintain an average amplitude of actuation when switching the open-loop actuation frequency or duty cycle, or when switching between closed-loop control laws. We will discuss some effects of the different settings of gain of the actuation amplitude controller in section 4.7.

An average amplitude of actuation can be represented in a non-dimensional form as the momentum coefficient:

$$C_\mu = \frac{\dot{m}_a}{\dot{m}_\theta}, \quad (2.1)$$

where \dot{m}_a is the average mass flow rate of compressed air used by the actuator system. It is measured at the inlet of the plenum chamber by a *Brooks* 5863S mass flow meter. An estimation of the mean mass flow rate through the upper boundary layer is defined as $\dot{m}_\theta = \rho \theta l_z U_1$. Here, θ is the momentum thickness of the boundary layer, ρ is the density of air, l_z is the splitter plate span, and U_1 is the mean velocity of the upper stream.

The target amplitude of actuation is imposed by setting the plenum pressure to a desired level. Since the compressed air system cannot be controlled in-time, the instantaneous amplitude depends on the frequency of actuation (f_a) and the duty cycle (dc). Figure 3 shows the dependence of C_μ , with respect to the available range of open-loop actuation settings. Based on these results we can approximate a constant average amplitude of actuation for $f_a \leq 50$ Hz, for a given constant duty cycle. These results are obtained for a plenum pressure level $P_t = 15$ mbar. This pressure level is used throughout this study except where stated otherwise. It corresponds to an average nozzle velocity of the jets of around 3 m s^{-1} , which is very close to the convective speed of the mixing layer. Additional information on the velocity response of the jet actuators is available in Parezanovic *et al.* (2014).

2.3. Sensor system

The main state evaluation and control feed-back sensor system is a rake of up to 24 hot-wire probes. The hot-wire probes are of a classic single wire design, capable of measuring a modulus of velocity in a single plane (xy plane in our case). In a mixing layer, the cross-stream component of velocity is one order of magnitude lower than the stream-




FIGURE 3. Actuator system output C_μ for actuation frequencies $1 < f_a < 400$ Hz, and duty cycles (dc) of 15 %, 30 %, 50 % and 70 %, using a plenum pressure setting of $P_t = 15$ mbar. Actuation frequencies f_a are plotted on a logarithmic scale, on the horizontal axis. Momentum coefficient C_μ , plotted on the vertical axis, is computed for the low speed mixing layer conditions (see section 4.1).

wise component. Therefore, the contribution of the cross-stream component of velocity fluctuations v' can be neglected, and the hot-wire sensors can be considered sensitive only to the stream-wise component u' .

The rake can be positioned at key points of interest in the mixing layer flow by a 2D displacement console (in xy plane) as illustrated in figure 1. The sensors cover a vertical length of up to 184 mm depending on the number of wires used and their placement on the rake. Anemometer units are of in-house design using TSI 1750 constant temperature anemometry modules. They are calibrated for an optimal signal response up to 20 kHz. The hot-wire velocity measurements are sampled at a rate of 5 kHz, with a precision of 0.05 m s^{-1} , and are corrected for temperature drifts using a reference temperature sensor at the inlet of the test section.

Two *Lavision Image Intense* cameras coupled with a *Spectra-Physics Lab-130* Nd:YAG laser are used for flow visualization in xy plane. Seeding for visualization is provided by vaporized oil particles.

2.4. Real-time I/O system

Acquisition of data and real-time signal processing for control are performed by a *Concurrent iHawk* Real-Time computing system. This system enables simultaneous data acquisition of up to 64 analog input channels and signal output to 96 digital output channels. The analog input channels are managed by two acquisition cards each with 32 channels. In the current experimental setup, the sensor signals are duplicated so that both cards acquire the same data from the hot-wire rake. One of the cards is dedicated to acquiring and relaying hot-wire signals as controller sensor input at a rate of 1 kHz, while the second card is recording the data for state evaluation and post-processing with a higher resolution of 5 kHz. The system uses in-house developed software for data recording, and Concurrent's Simulation Workbench software package for real-time I/O and processing operations.

3. Closed-loop control strategy

In this section we define the objective functions used for evaluation of both open- and closed-loop control. Next, we introduce Machine Learning Control (MLC) method and the basic principles of Genetic Programming (GP), on which this method is based. We then discuss specific aspects of implementation of MLC in the mixing layer experiment.

3.1. Objective functions

The primary goal of the experiment is manipulation of the local flow properties of a turbulent mixing layer. The hot-wire sensors allow us to recover a local velocity fluctuation profile at the point of interest in the mixing layer. The total velocity u_i measured by the i -th hot-wire is then decomposed in a sum of its mean and fluctuating parts, leading to:

$$u'_i(t) = u_i(t) - \langle u_i(t) \rangle_\tau \quad (3.1)$$

where

$$\langle u_i(t) \rangle_\tau = \frac{1}{\tau} \int_{t-\tau}^t u_i(t) dt \quad (3.2)$$

is the time-averaged mean velocity estimated on a period of $\tau = 2$ s, which corresponds to around 200 Kelvin-Helmholtz vortices.

Based on these sensor inputs, we propose two different objective functions:

$$K = \sum_{i=1}^{N_s} \langle u_i'^2(t) \rangle_T, \quad (3.3)$$

and

$$W = \frac{2K}{\max_{i \in [1, N_s]} \langle u_i'^2(t) \rangle_T}. \quad (3.4)$$

In (3.3) and (3.4), N_s is the number of hot-wire probes used in the experiment (see section 3.3) for estimating the objective functions, and $T = 10$ s is the averaging time for obtaining velocity variances $\langle u_i'^2 \rangle_T$. This value corresponds to a typical evaluation time of open-loop actuation or a closed-loop control law candidate. The objective function K is calculated as the total sum of variances from all hot-wire sensors used. It is proportional to an integral 1D estimation of the turbulent kinetic energy in the mixing layer. The objective function W is based on the local thickness (or width, hence W) of the mixing layer inferred from the velocity fluctuation profile shape. It is designed to favor small gradients in the distribution of fluctuation rates along the profile, while still seeking a large total sum of variances. This is brought on by assuming that the best mixing would be represented by a homogeneous fluctuation profile. W is simply calculated by penalizing K with the maximum value of variances detected in the profile. If $\max_i \langle u_i'^2 \rangle_T$ is the height of a triangle which approximates the profile surface, then W is proportional to the base of that triangle, or an effective local thickness of the mixing layer. This is illustrated in figure 4 for two examples of velocity variance profiles encountered in the mixing layer.

For measuring the control effectiveness of the machine learning algorithm, we define non-dimensional cost functions based on K and W *i.e.*

$$J_K = \frac{K}{K_u} \quad \text{and} \quad J_W = \frac{W}{W_u}, \quad (3.5)$$

where K_u and W_u are reference values of the objective functions for the un-actuated

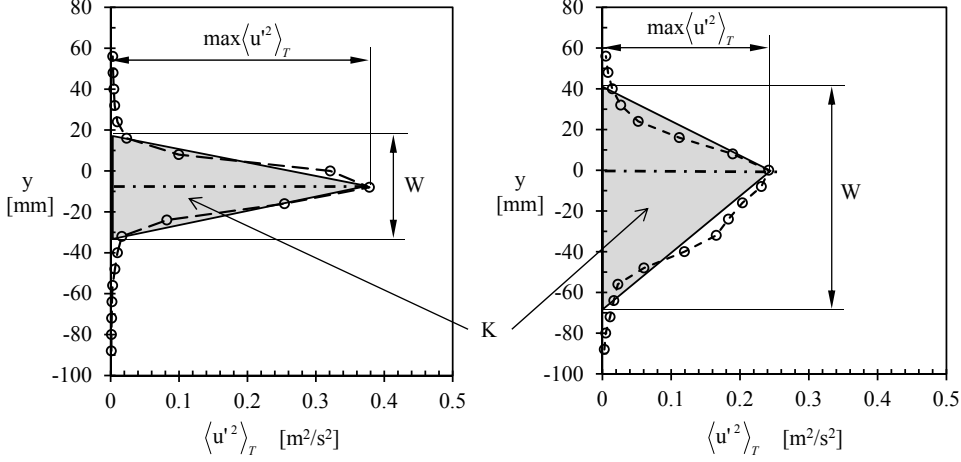


FIGURE 4. Physical interpretation of the objective functions K and W . Illustration on two different shapes of velocity variance profiles.

mixing layer. These cost functions correspond to a relative change, brought on by actuation, with respect to the un-actuated flow. Depending on the goal (see section 4), MLC seeks to maximize or minimize these cost functions, thereby enhancing or reducing the corresponding objective function, respectively.

These objective functions do not include any penalization terms regarding actuation cost. This choice was made taking into account the implementation of the PI control (see section 2.2) for maintaining an average constant actuation amplitude. In this case, instantaneous fluctuations of the actuation amplitude are acceptable, and can allow MLC to potentially evolve solutions qualitatively different from periodic or near-periodic forcing at a single frequency (see section 4.7). The energy efficiency of a control is evaluated, *a posteriori*, by considering the momentum coefficient C_μ given by (2.1).

3.2. Machine learning control

We explore the possibilities of designing a closed-loop control strategy for strongly non-linear problems in fluid dynamics in a model-free manner. We are particularly interested in developing a sensor feedback approach where the control, to be determined, is a function of some sensor signals. This problem is equivalent to finding the regression model f_{MLC} such that

$$b(t) = f_{\text{MLC}} \left(s_{i \in [1, N_s]}(t) \right) \quad (3.6)$$

where b is the control law, and s_i are the sensor signals. For determining f_{MLC} , we turn to the interdisciplinary field of machine learning (Murphy 2012) and in particular to one of its well-established techniques for symbolic regression, called genetic programming (Koza 1992). The advantage of symbolic regression over standard regression methods is that in symbolic regression, the search process works simultaneously on both the model specification problem and the problem of fitting coefficients. The genetic programming paradigm has been successfully employed to solve a large number of difficult problems such as pattern recognition, robotic control, theorem proving and air traffic control. To the authors' best knowledge, genetic programming has never been implemented in control of fluid flows by other authors.

Genetic programming (GP) is a biologically-inspired algorithm introduced by Koza to

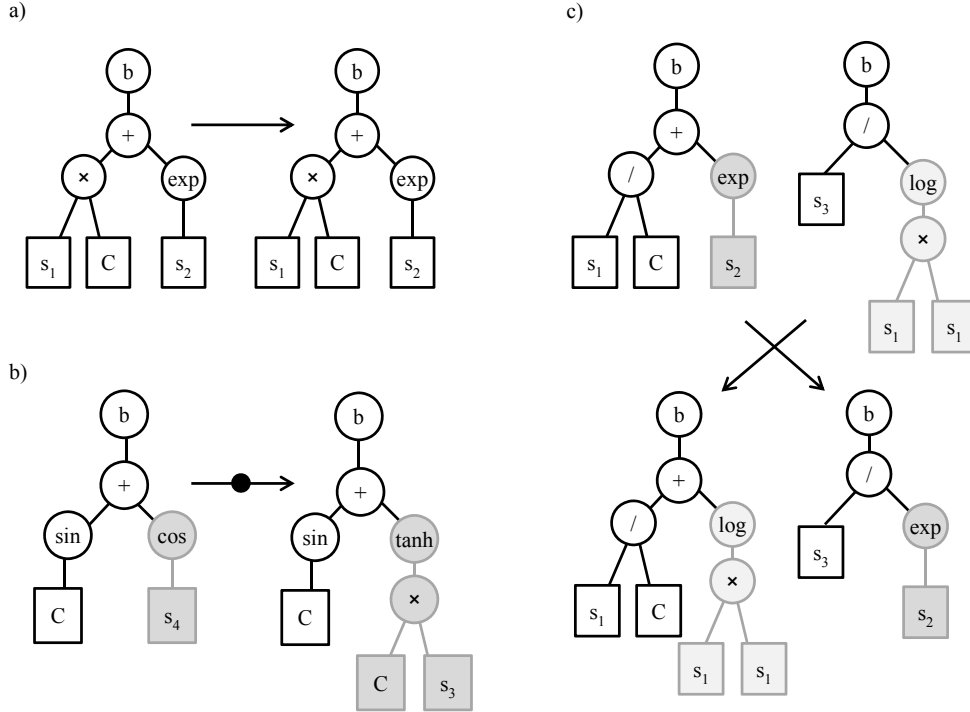


FIGURE 5. Biologically-inspired operations performed by Genetic Programming (GP) on the tree-like functions: a) replication, b) mutation and c) cross-over. The tree represented in a) corresponds to the function $b(t) = C \times s_1(t) + \exp(s_2(t))$ where C is a constant to be determined by GP. For the mutation b), the operation consists of selecting a node, erasing its subtree and replacing it by another one created randomly. Part of the information contained in the individual is kept while new information is allowed to enter the population. Mutation increases the diversity and is responsible for exploring the search space with large steps. For the cross-over c), one node in each of the two individuals selected for the cross-over is randomly chosen, the nodes and their subtrees are then exchanged. The cross-over is responsible for exploring the search space around well-performing individuals.

find computer programs that perform a user-defined task. It constitutes an extension of the genetic algorithm (Goldberg 1989) for which the genetic operators (selection, cross-over and mutation) operate on a set of operations, elementary functions, variables and constants. When GP is used for symbolic regression, it combines automatically these elements to search for a symbolic expression that constitutes the best solution to a given optimization problem. Proposed symbolic expressions (examples shown in figure 5) are built as tree-like structures (Koza 1992) which can be easily evaluated in a recursive manner and described by a LISP expression. These trees are called individuals, while a population of individuals is a generation. GP is implemented as an iterative procedure in which a population of candidate solutions evolves toward better solutions by repeatedly undergoing genetic modifications based on their fitness.

The first generation of individuals is created in a random manner. All individuals are evaluated and a fitness value is attributed based on how well they minimize or maximize the objective function. Based on these evaluations, a second generation of individuals is created using three principal genetic operations: replication, mutation and cross-over. Each of the genetic operations occurs with a pre-determined probability. If replication

is selected then the candidate individual will simply be copied into the next generation (figure 5a). Mutation causes a part of the individual's tree to be replaced by a newly created random sub-tree as shown in figure 5(b). A cross-over operation involves a pair of individuals, which will exchange a randomly selected sub-tree between each other and two new individuals created in such a manner will become a part of the next generation (figure 5c).

The selection process of candidate individuals uses the tournament method; few random individuals are chosen to compete in a tournament and the winner (based on its evaluated fitness) is selected for one of the genetic operations to be performed on it. The tournament is an efficient way of controlling the selectiveness of the GP evolutionary process. Usually, the size of the tournament is small compared to the size of the population. The bigger the tournament size N_t is, the more probable it is to include the overall best individuals as participants, leading to a more selective process. The selection ranges from a random individual for $N_t = 1$, to selection of the best individual of the generation if N_t equals the size of the population of the current generation.

An additional operation available is elitism, which copies the single or few best individuals of one generation directly into the next generation, while avoiding the tournament process. This operation is generally used to ensure that the best individuals of one generation are not lost and remain available for further improvements in the future generations.

When all the individuals of the next generation are created, they are evaluated and the selection process begins anew in order to build a subsequent generation. This iterative process continues for a desired number of generations. A rule of thumb is that given a sufficient number of individuals in each generation, a solution should be obtained in less than $N_g = 50$ generations. There is no mathematical proof of convergence, but the method has been successful in many applications (Lewis *et al.* 1992; Nordin & Banzhaf 1997).

3.3. MLC in the mixing layer experiment

The design layout of the experimental machine learning control is presented in figure 6. We can arbitrarily separate the system into three parts for analysis. The top part is the experiment with the mixing layer plant and the associated actuator and sensor systems. This part represents the un-actuated flow and the elements sufficient for open-loop control; a function generator can be used to drive the micro-jets directly.

The middle part is an interpreter/controller system which operates in real-time (cyclic rate of 1 kHz); sensor information is continuously acquired and evaluated through a control law to produce the control signal as in (3.6). An Heaviside function with a threshold at 0 is applied to the evaluated signal in order to produce a binary command, which triggers the jet actuators. For minimum operation, this part requires an expression representing a control law (for example $b(t) = 0$, to switch off actuation). This expression can be input manually, but the principal way of obtaining a control law expression is through MLC.

The bottom part, shown in figure 6, represents the learning phase. MLC sends a candidate individual to the real-time interpreter for testing. During the evaluation, data from the sensors are recorded simultaneously with data used for feedback. At the end of the evaluation time, a fitness value is calculated and assigned to the tested individual. This information is stored in the MLC database and the next candidate individual is sent for evaluation. The procedure is repeated until the final generation, from which the best individual is selected as the best closed-loop control law. This individual can then be applied continuously in the interpreter/controller and the learning phase is disconnected.

The genetic programming part of MLC code is developed in-house based on an open-

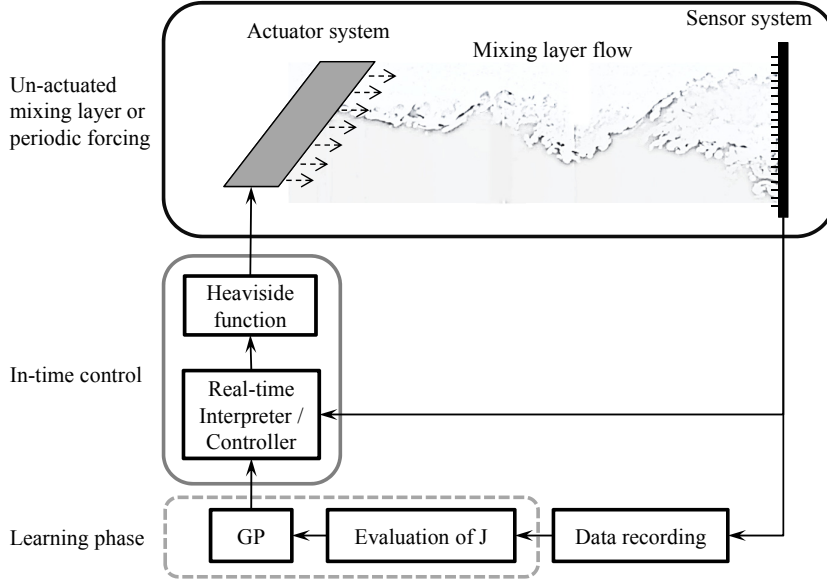


FIGURE 6. The layout of MLC: the mixing layer plant with sensors and actuators, the in-time interpreter/controller and the MLC learning phase.

source code *ECJ* (Luke *et al.* 2013). Genetic programming codes using numerical simulations as evaluation data do not need to evaluate the same individual more than once. In the experiment, however, one is faced with uncertainties and stochasticity created by natural perturbations in the flow and the actuator system, measurement precision, etc. In order to be better adapted to experimental conditions, our MLC code contains crucial new features:

- (a) An individual is evaluated every time it appears in the genetic programming process.
- (b) The cost of an individual is averaged with the values recorded previously for the same individual, resulting in a cumulative moving average being used as the cost function value.
- (c) A predetermined number of best individuals in each generation are re-evaluated several more times to ensure a stable sample size for the averaging process.
- (d) No duplicate individuals are allowed during the creation of the first generation.

The GP learning module thus assigns a cumulative average cost function value (fitness) to an individual. The objective is to choose as a best individual, one which has a continuously good performance, rather than one with strongly varying performance due to external perturbations. Such perturbations may cause an undeserved high fitness value to be attributed to an otherwise unremarkable individual, which is the only real danger to the convergence process. The averaged cost approach is designed to promote robustness and prevent such events from having an influence, but at a price of being much more conservative in the selection of individuals. If a duplicate individual is detected during the creation of the first generation, it is rejected and replaced with a new random individual. This ensures that the search space is not reduced by probing the same point two or more times.

Operations and elementary functions used in the experiment are: $+$, $-$, \times , $/$, \sin , \cos ,

Parameter	Value
Number of generations	$N_g = 25$
Population size	$N^{(1)} = 1000,$ $N^{(n)} = 100 \ (n = 2, \dots, 25)$
Tournament size	$N_t = 7$
Replication	$P_r = 0.1$
Cross-over	$P_x = 0.65$
Mutation	$P_m = 0.25$
Elitism	$N_e = 1$
Min. tree depth	3
Max. tree depth	10
Operations	$+, -, \times, /, \sin, \cos, \exp, \tanh, \log$

TABLE 1. Standard set of MLC parameters used in the experiments (except where noted otherwise).

\exp , \tanh and \log . The variables available to MLC are instantaneous values of velocity fluctuations from the hot-wire sensors *i.e.* $s_i(t) \equiv u'_i(t)$ in (3.6). Every third sensor in the rake is used as an MLC variable since more would only carry redundant information for the sensor-based controller. Hence, $N_s = 21$ for the evaluation of the cost functions (3.3) and (3.4), while only 7 sensors across the shear layer are employed in (3.6). Finally, constants used are in a range of $[-1, 1]$ with a precision up to a second decimal. The range and precision are adapted to be of the same order of magnitude as the typical velocity fluctuation information given by the sensors.

In the experiment, the standard rates for the genetic programming operations are 10 % for replication, 25 % for mutation and 65 % for cross-over. Elitism is set to $N_e = 1$, meaning a single best individual in a generation is sent directly to the next generation. An experiment consists of $N_g = 25$ generations, where the first generation is composed of $N^{(1)} = 1000$ individuals to allow for increased initial diversity, while subsequent generations ($n = 2, \dots, 25$) contain $N^{(n)} = 100$ individuals each. An individual is limited to a tree depth of 10 levels. Standard MLC parameters are given in table 1. Some variations of these settings are explored in section 5.

The evaluation time of every individual in the experiment is $T = 10$ s, as mentioned earlier in this section. In addition, there is up to 6 s of various delays in order to facilitate communication between the learning module, the controller and the data recording system. Also, a few seconds must be allowed for the mixing layer to register the effects of each new actuation when changing individuals during the learning process. Taking this into account, a single evaluation of an individual takes around 20 s to complete. A standard MLC experiment (see table 1) is completed in approximately 26 hours.

At the start of each generation, the un-actuated flow and the best open-loop actuation are re-tested in order to keep track of these reference values. Hence, each individual's cost (J_K or J_W) is calculated using the reference baseline flow values (K_u or W_u) obtained at the beginning of its own generation.

4. Results

In this section, we first present the initial conditions of the un-actuated mixing layer, for the Low-Speed (LS) and High-Speed (HS) configurations. This is followed by a detailed mapping of the impact of open-loop forcing on the cost functions J_K and J_W . We will then focus on the best closed-loop control laws found by MLC for several different cases. Maximization of cost functions J_K (MaxK) and J_W (MaxW) is explored at $x = 200$ mm and $x = 500$ mm, as well as minimization of J_K (MinK) at $x = 200$ mm. The impact of different flow conditions is explored with a maximization experiment of J_W for the HS mixing layer configuration, at $x = 200$ mm. Lastly, the effect of alternative gain settings of the actuation amplitude controller is tested in an MLC experiment. These settings cause the amplitude controller to become affected by actuation, and MLC is capable of exploiting this coupling to produce a different type of a solution. In sections dealing with MLC results, values for the un-actuated flow, best open-loop reference and the best individual are evaluations performed in the final generation of each experiment. Therefore, they represent results of single evaluations of the mentioned configurations. Statistically converged results of these cases will be discussed in section 5.

4.1. Un-actuated flow

Two configurations of the mixing layer are used in the experiment: a low-speed (LS) mixing layer with stream velocities of $U_1 = 4.7 \text{ m s}^{-1}$ and $U_2 = 1.3 \text{ m s}^{-1}$, and a high-speed (HS) configuration with stream velocities of $U_1 = 9 \text{ m s}^{-1}$ and $U_2 = 1.7 \text{ m s}^{-1}$. In each configuration, the magnitude of velocity U_1 is selected to have a stable boundary layer on the upper surface of the splitter plate, laminar in the LS configuration, and turbulent for the HS setting.

The two streams of the LS configuration produce a mixing layer with a velocity ratio of $r = 0.27$, and a convective velocity of $U_c = 3 \text{ m s}^{-1}$. The boundary layer on the high speed side of the splitter plate is laminar ($\delta_{99} = 6.9 \text{ mm}$, $\theta = 0.73 \text{ mm}$ and $H = 2.44$). Based on the momentum thickness of the upper boundary layer, the initial Reynolds number of the mixing layer can be estimated as $Re_\theta = 500$.

The HS configuration features a turbulent upper boundary layer ($\delta_{99} = 13.3 \text{ mm}$, $\theta = 0.81 \text{ mm}$ and $H = 1.55$). Initial conditions for this configuration are $Re_\theta = 2000$, $U_c = 5.4 \text{ m s}^{-1}$ and $r = 0.18$.

In both mixing layer configurations, the thickness of the boundary layer on the lower side of the splitter plate is $\delta_{99} < 1 \text{ mm}$. This is a result of very low velocity U_2 (in both cases) and the inserted foam flow stabilizer. The impact of this boundary layer on the development of the mixing layer can be considered negligible. The main focus in this section will be on the low-speed configuration. The high-speed configuration will be examined in section 4.6, and later discussed with respect to control robustness in section 5.3.

Based on the flow visualization of the low-speed mixing layer, shown in figure 7, we can estimate the frequency of the initial Kelvin-Helmholtz instability as $f_{KH} = 90 \text{ Hz}$. Two streamwise locations are shown in figure 7, one at $x = 200 \text{ mm}$ and the second at $x = 500 \text{ mm}$, with their respective mean velocity and velocity variance profiles. These locations are used throughout this study and are selected as points of interest based on mixing layer spatial evolution data presented in more detail in Parezanovic *et al.* (2014).

4.2. Periodic forcing

The response of the mixing layer to periodic forcing is evaluated using objective functions J_K and J_W on the entire frequency range $1 < f_a < 400 \text{ Hz}$ available to the actuator system. Measurements are performed for the sensor system placement at $x = 200 \text{ mm}$

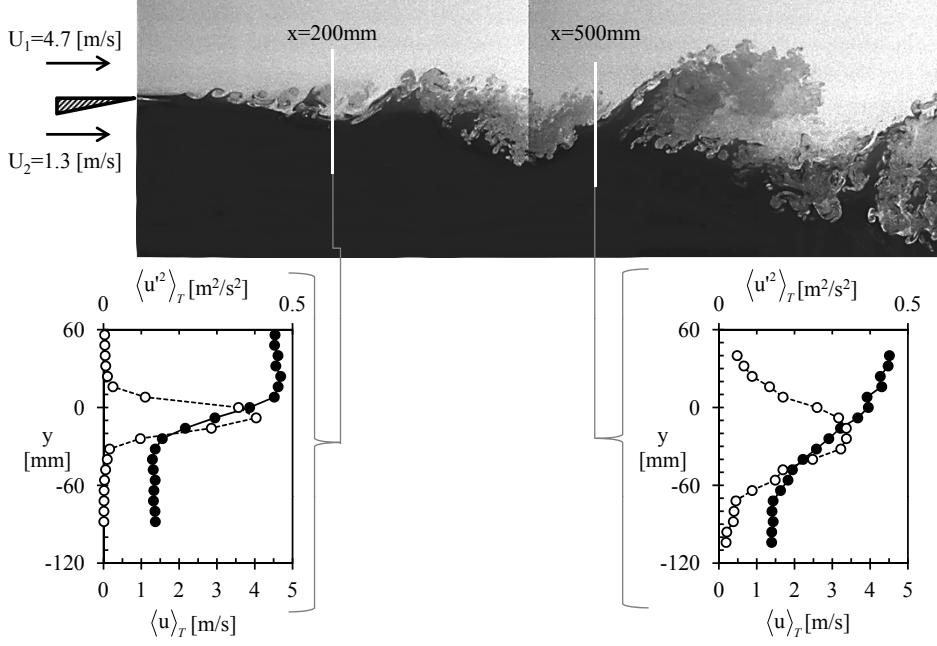


FIGURE 7. Smoke visualization of the un-actuated low-speed mixing layer with mean velocity $\langle u \rangle_T$ and velocity variance $\langle u'^2 \rangle_T$ profiles for sensor positions at $x = 200$ mm and $x = 500$ mm.

and $x = 500$ mm. Actuation frequency is changed in varying steps (roughly a logarithmic distribution) using 60 points in the 400 Hz range. Four different actuation duty cycles are used as mapping parameters: 15 %, 30 %, 50 % and 70 %.

Figure 8 shows the evolution of the cost functions J_K and J_W with regard to open-loop actuation frequency when sensors are placed at $x = 200$ mm and $x = 500$ mm. The values for the un-actuated case are equal to 1. For the sensor location at $x = 200$ mm (left in figure 8), the most striking increase of J_W can be observed for $dc_{W_{\max}}^{\text{OL}} = 50$ % in a very narrow band of actuation frequencies around $f_{W_{\max}}^{\text{OL}} = 21$ Hz. This result corresponds to a thick velocity variance profile in the associated diagram. Values of J_K are increased the most for $f_{K_{\max}}^{\text{OL}} = 12$ Hz and $dc_{K_{\max}}^{\text{OL}} = 70$ %. For higher actuation frequencies of $f_a > 75$ Hz the values of J_K are reduced below the value of the un-actuated flow in all four cases. This result corresponds to high frequency forcing effect on the near-field of a free shear layer, reported by Vukasinovic *et al.* (2010). However, a relation $J_W > 1$ stands for the whole range of tested actuation frequency and duty cycle combinations.

Actuation effects measured further downstream at $x = 500$ mm are similarly shown in figure 8(right). A highest gain with regard to J_W is again obtained for $dc_{W_{\max}}^{\text{OL}} = 50$ %, but this time it is at a lower frequency $f_{W_{\max}}^{\text{OL}} = 9$ Hz. Evolution of J_K indicates two positive extremum points, with the highest gain on the lower frequency extremum with the actuation parameters $f_{K_{\max}}^{\text{OL}} = 5$ Hz, $dc_{K_{\max}}^{\text{OL}} = 70$ %. The minimum for J_K is obtained around $f_{K_{\min}}^{\text{OL}} = 75$ Hz for $dc_{K_{\min}}^{\text{OL}} = 30$ % and 50 %, although it is not as pronounced as for the measurements at $x = 200$ mm. W is always larger than the un-actuated flow value in this case as well. These optimal forcing parameters will be used in the MLC experiments as best actuation settings f_a^{opt} and dc^{opt} for reference open-loop control values, and are given in table 2.

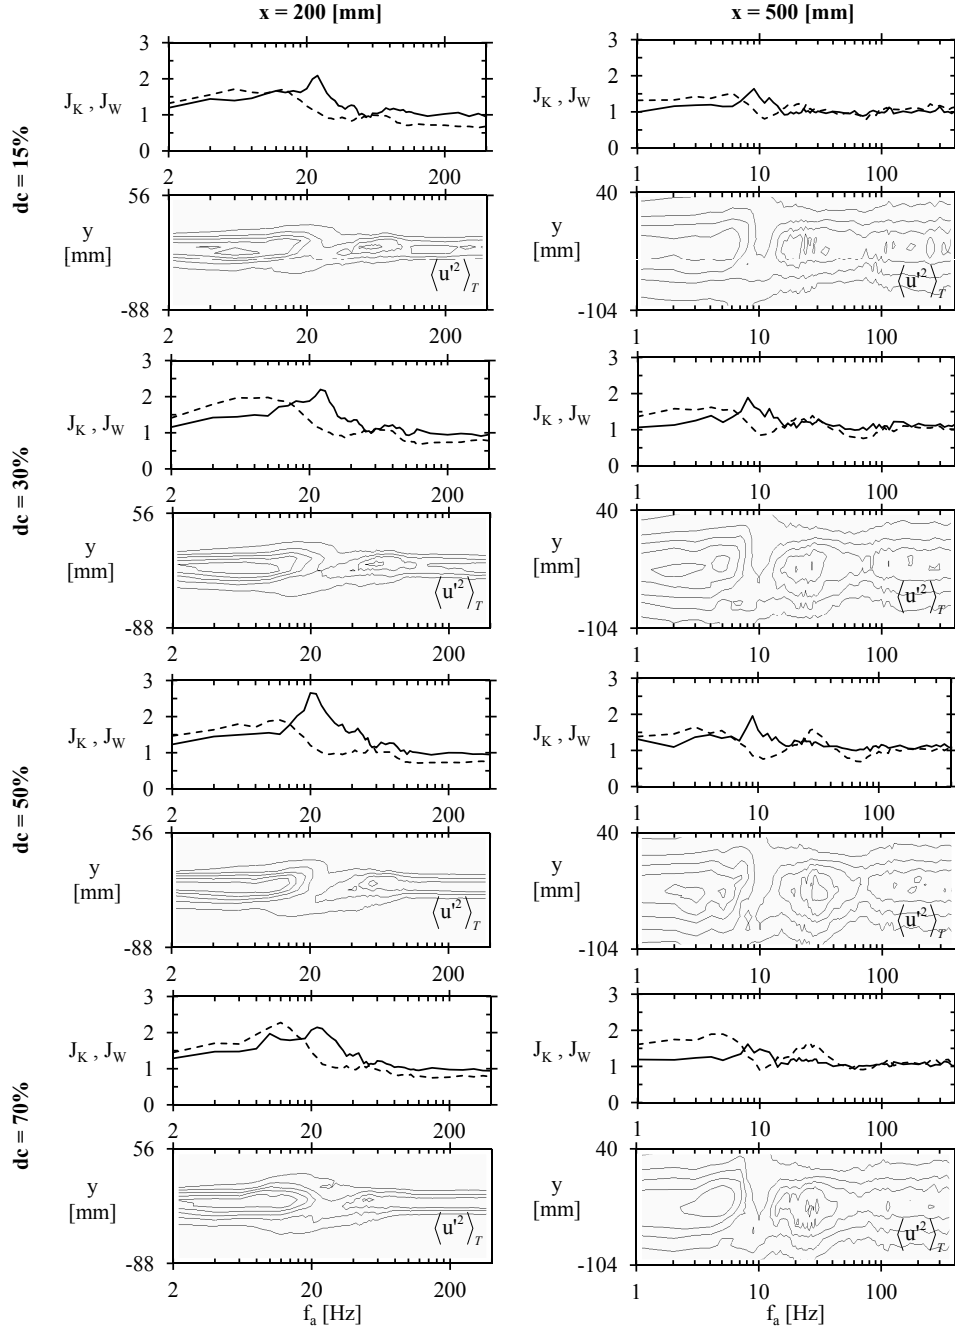


FIGURE 8. Mapping of mixing layer response to open-loop actuation for $f_a = 0 - 400$ Hz at $x = 200$ mm (left) and $x = 500$ mm (right). Duty cycle settings (top to bottom): $dc = 15\%$, $dc = 30\%$, $dc = 50\%$ and $dc = 70\%$. All diagrams share a common horizontal axis denoting f_a on a logarithmic scale. Cost functions are represented with continuous line for J_W and dashed line for J_K . Corresponding mean velocity variance profiles for each case are plotted as iso-lines of $\langle u'^2 \rangle_T$, with five levels in range $[0.1 \ 0.5] \text{ (m}^2 \text{ s}^{-2})$. Vertical axis for profile plots corresponds to the vertical position of the sensors.

4.3. MLC maximization of the width (MaxW experiment)

MLC control is first applied to search for a maximization of the objective function J_W at $x = 200$ mm and $x = 500$ mm. These optimization problems are denoted MaxW. The GP parameters are identical in both cases and correspond to standard parameters given in table 1. The best performing individuals for the two experiments are selected in the last MLC generation *i.e.* $n = 25$.

Figure 9.1(a) shows velocity variance profiles for the un-actuated flow, best open-loop actuation and the best MLC individual. We can note that the fluctuation profiles of the two actuated cases are almost indistinguishable from each other but significantly different from the un-actuated flow profile. Figure 9.1(b) shows the FFT spectra of the actuation signals for the open- and closed-loop control. This allows us to determine a dominant actuation frequency $f_{W_{\max}}^{\text{MLC}}$ for the best closed-loop individual and compare with the best (fixed) open-loop actuation frequency $f_{W_{\max}}^{\text{OL}}$. It can be noted that in this experiment the best closed-loop had a dominant frequency $f_{W_{\max}}^{\text{MLC}} = 19$ Hz slightly lower than the best open-loop frequency $f_{W_{\max}}^{\text{OL}} = 21$ Hz. Figure 9.1(c) shows pseudo-visualization of the respective flows, based on Taylor's hypothesis of frozen turbulence, using the velocity fluctuation u' time-series from hot-wire probes. The time series represented corresponds to 1 s which is equivalent to a spatial length of around 10 times of the sensor rake height used in this experiment. It can be concluded that the best MLC individual performs very similarly to the best open-loop control. They both increase the local thickness of the mixing layer, as evidenced from the velocity variance profiles. The visualization reveals very similar sizes and frequencies of coherent structures.

Results are displayed in a similar fashion in figure 9.2 for the $x = 500$ mm case. This time, there are some differences in the velocity variance profiles (figure 9.2(a)) between the controlled cases, but it can be said that both work to create a similar mixing layer thickness increase. Spectra in figure 9.2(b) show the MLC actuation frequency is identical to open-loop setting: $f_{W_{\max}}^{\text{MLC}} = f_{W_{\max}}^{\text{OL}} = 9$ Hz. However, from the actuation stripe on top of pseudo-visualization fields, we can observe a much different instantaneous duty cycle for the closed-loop case.

Table 2 provides quantitative information to compare the optimal parameters corresponding to MaxW for the two measurement locations. Actuation frequency and duty cycle are fixed quantities in the case of open-loop actuation. For the closed-loop case, the reported actuation frequency is the dominant frequency in the spectrum of the best closed-loop actuation signal *i.e.* $f_{W_{\max}}^{\text{MLC}}$. The duty cycle reported for this case is estimated as the percentage of total evaluation time for which the actuators are in the active state. The momentum coefficient C_μ represents the invested energy of actuation. Finally, the values for the cost function J_W are given.

It can be concluded that the best closed-loop individuals can outperform their open-loop counterparts in all aspects. In particular, an interesting case is the optimal duty cycle for closed-loop control at $x = 500$ mm which is selected around 25 %. This corresponds well in comparison with the open-loop response maps in section 4.2, where figure 8 shows that the maximum of J_W at $x = 500$ mm is similar for $dc = 30$ % and $dc = 50$ %. Duty cycle $dc = 25$ % was not tested in open-loop mapping experiments, but the best MLC individual points in the direction of that result. However things may not be that simple since in the case of $x = 200$ mm, MLC has found repeatedly that the best resulting closed-loop actuation has a dominant frequency of 19 Hz with an optimal duty cycle of 48 %. This duty cycle is very close to the optimal open-loop value of 50 %, yet the frequency for open-loop should not be optimal below at least 20 Hz.

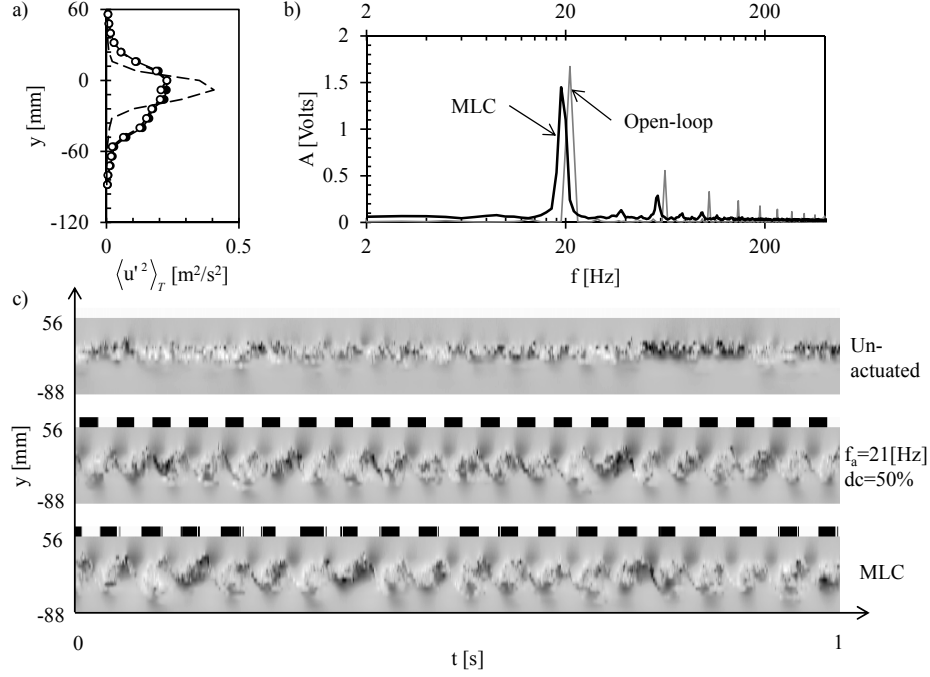

9.1) $x = 200$ mm.9.2) $x = 500$ mm.

FIGURE 9. Results of MaxW experiments: a) profiles of $\langle u'^2(t) \rangle_T$ for the un-actuated flow (dashed line), open-loop reference (white circles) and MLC (black circles), b) actuation signal spectra for open-loop (grey) and MLC (black) and c) pseudo-visualization using hot-wire time series of u' (gray-scale) in range $[-1.5, 1.5]$ (m s⁻¹) with a stripe depicting actuation activity (black stands for actuation).

Maximization of W (MaxW)				
$x = 200$ [mm]	f_a^{opt} [Hz]	dc^{opt} [%]	C_μ	J_W
Open-loop (OL)	21	50	0.25	2.13
Closed-loop (MLC)	19	48	0.25	2.25
$x = 500$ [mm]	f_a^{opt} [Hz]	dc^{opt} [%]	C_μ	J_W
Open-loop (OL)	9	50	0.20	1.54
Closed-loop (MLC)	9	25	0.13	1.70
Maximization of K (MaxK)				
$x = 200$ [mm]	f_a^{opt} [Hz]	dc^{opt} [%]	C_μ	J_K
Open-loop (OL)	12	70	0.27	2.27
Closed-loop (MLC)	12	62	0.24	2.44
$x = 500$ [mm]	f_a^{opt} [Hz]	dc^{opt} [%]	C_μ	J_K
Open-loop (OL)	5	70	0.25	1.98
Closed-loop (MLC)	4.5	80	0.34	1.95

TABLE 2. Results of the final generation of the MLC maximization experiments of W (MaxW) and K (MaxK) for two positions of the sensors $x = 200$ mm and $x = 500$ mm. The optimal values of f_a^{opt} and dc^{opt} used for the open-loop reference were determined in section 4.2.

4.4. MLC maximization of the fluctuation energy (MaxK experiment)


Results for the maximization of the energy-based objective function K are presented here for $x = 200$ mm and $x = 500$ mm. Resulting best individuals are selected from generation $n = 25$ from standard MLC experimental runs.

Mean velocity variance profiles shown in figure 10.1(a) for the results obtained at $x = 200$ mm, reveal a slight advantage for the closed-loop control, while the global frequency for closed-loop control matches the open-loop value, shown in figure 10.1(b). Visualization fields in figure 10.1(c) indicate no significant differences between the two controls, at first glance. Data in table 2 under the $x = 200$ mm entry for MaxK show a slight optimization of the closed-loop duty cycle which could be responsible for improved performance.


Results are similarly presented in figure 10.2 for maximization of K at $x = 500$ mm. Here (figure 10.2a) we note larger differences in the velocity variance profile shapes between the controlled cases. Actuation frequency is slightly lower for closed-loop (shown in figure 10.2b). Pseudo-visualization in figure 10.2(c) reveals a large increase of the actuation duty cycle for the case of the closed-loop control. Results for MaxK at $x = 500$ mm, in table 2, show that open-loop is slightly better and much more efficient in terms of invested energy.

4.5. MLC minimization of the fluctuation energy (MinK experiment)

In section 4.2 it has been shown that the total energy of fluctuations in the mixing layer could be reduced using high frequency periodic actuation. An MLC experimental run is performed in search of the best K minimization closed-loop control and results are shown in figure 11.1. The velocity variance profiles in figure 11.1(a) show that both control approaches have a much smaller relative impact (compared to maximization experiments) when minimization of the mixing layer energy is attempted. The actuation frequency of the best MLC individual cannot be detected in figure 11.1(b). Figure 11.1(c)



10.1) $x = 200$ mm.



10.2) $x = 500$ mm.

FIGURE 10. Results of MaxK experiments (see caption of figure 9 for details).

reveals that continuous actuation was selected as the best closed-loop control. Resulting average reductions of K are 31 % for open-loop and 25 % for closed-loop control.

An MLC experiment seeking minimization of K has also been performed at $x = 500$ mm, but is not shown here. In this case, continuous actuation or no actuation are alternatively found as best individuals. This results from the fact that both have very similar objective function values at $x = 500$ mm, *i.e.* continuous actuation has some effect closer to the trailing edge but further downstream this effect is dispersed and the mixing layer instability re-appears. Note that high-frequency open-loop actuation still has the ability to reduce K at $x = 500$ mm as shown in section 4.2.

4.6. MLC maximization of the width in a high-speed mixing layer

The maximization of W has been explored in high-speed mixing layer flow conditions. The high-speed mixing layer is fully turbulent from its inception, however both low- and high-speed configurations are turbulent in nature at the sensor location at $x = 200$ mm downstream of the trailing edge. Still, the local frequency and velocity fluctuation levels are significantly different compared to the low-speed case. The mapping experiment of open-loop response for the high-speed configuration is performed and yields the best actuation parameters as $f_{W_{\max}}^{\text{OL}}(\text{HS}) = 35$ Hz and $dc_{W_{\max}}^{\text{OL}}(\text{HS}) = 50$ %, which are used as reference. Complete mapping results are omitted here for brevity.

Resulting velocity variance profiles and actuation frequencies, shown in figure 11.2(a) and (b), are very similar for both control types. Visualization in figure 11.2(c) reveals that closed-loop control is less regular than in previously documented cases for this measurement location, in low-speed conditions. This attests to the more turbulent nature of the high-speed mixing layer. The best MLC individual shows performance very close to best open-loop: $J_{W_{\max}}^{\text{OL}}(\text{HS}) = 1.72$ and $J_{W_{\max}}^{\text{MLC}}(\text{HS}) = 1.74$.


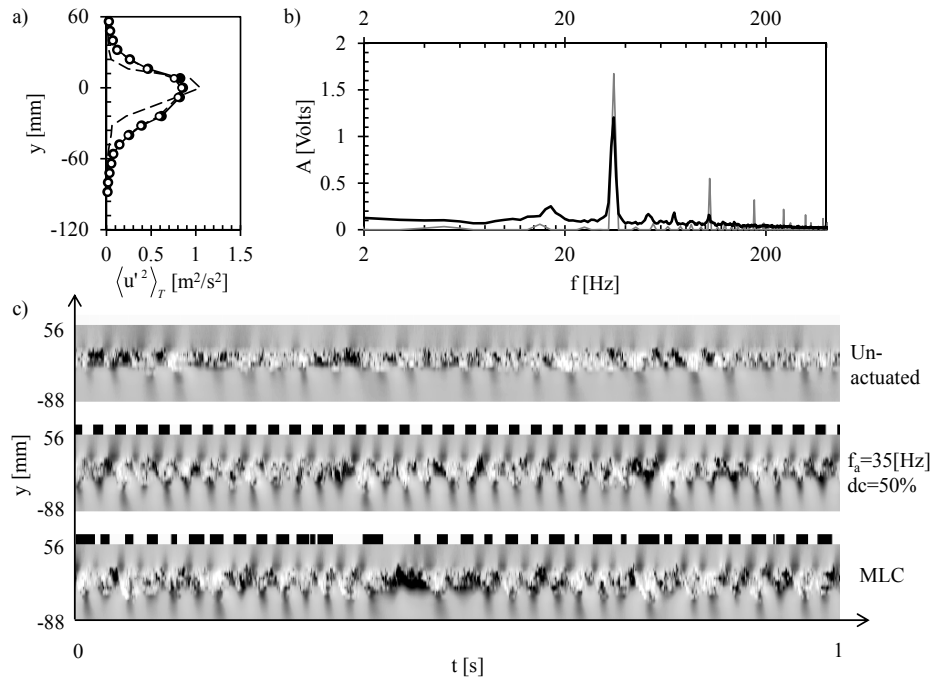
4.7. Influence of the high-gain amplitude controller on the MLC

Some properties of the PI controller, used for actuation amplitude control, were discussed in section 2.2. In the early experiments, a higher gain was used for the PI controller, making it very sensitive. Combined with the slow response of the compressed air distribution system, the effect was such that a low-frequency modulation introduced into the actuation signal caused the actuation amplitude to be significantly increased past the desired level on a short time scale, while preserving the long time scale target average amplitude.

Figure 12 shows two results from this experiment: (a) is the best individual from the first generation and (b) the best individual from generation $n = 50$. In each case the upper plot shows the instantaneous pressure levels in the plenum chamber of the actuator system, which corresponds to the pseudo-visualization time series shown in the lower plots. Except for the number of generations N_g , other parameters used in this experiment correspond to standard values in table 1.

The first generation individual (figure 12a) yields a similar value of J_K as the best open-loop actuation. However, due to the sensitive PI control, there exists a better solution (figure 12b), which employs bursts of actuation followed by comparatively long periods of continuous blowing. The action of continuous blowing drains the plenum chamber and the PI controller increases the flow rate to compensate. When the actuation burst starts, *i.e.* at the first short closure of the actuators, an over-pressure is created in the plenum. Each cycle in the burst is then injecting much more momentum into the mixing layer, causing a high instantaneous actuation amplitude, which yields a larger value of K . As the PI controller catches up with this process, another period of continuous blowing starts and the cycle repeats indefinitely.

This sequence comprises an actuation frequency of around 12 Hz (corresponding to the

11.1) Minimization of K (low-speed).11.2) Maximization of W (high-speed).FIGURE 11. Results of MinK (low-speed) and MaxW (high-speed) experiments at $x = 200$ mm (see caption of figure 9 for details).

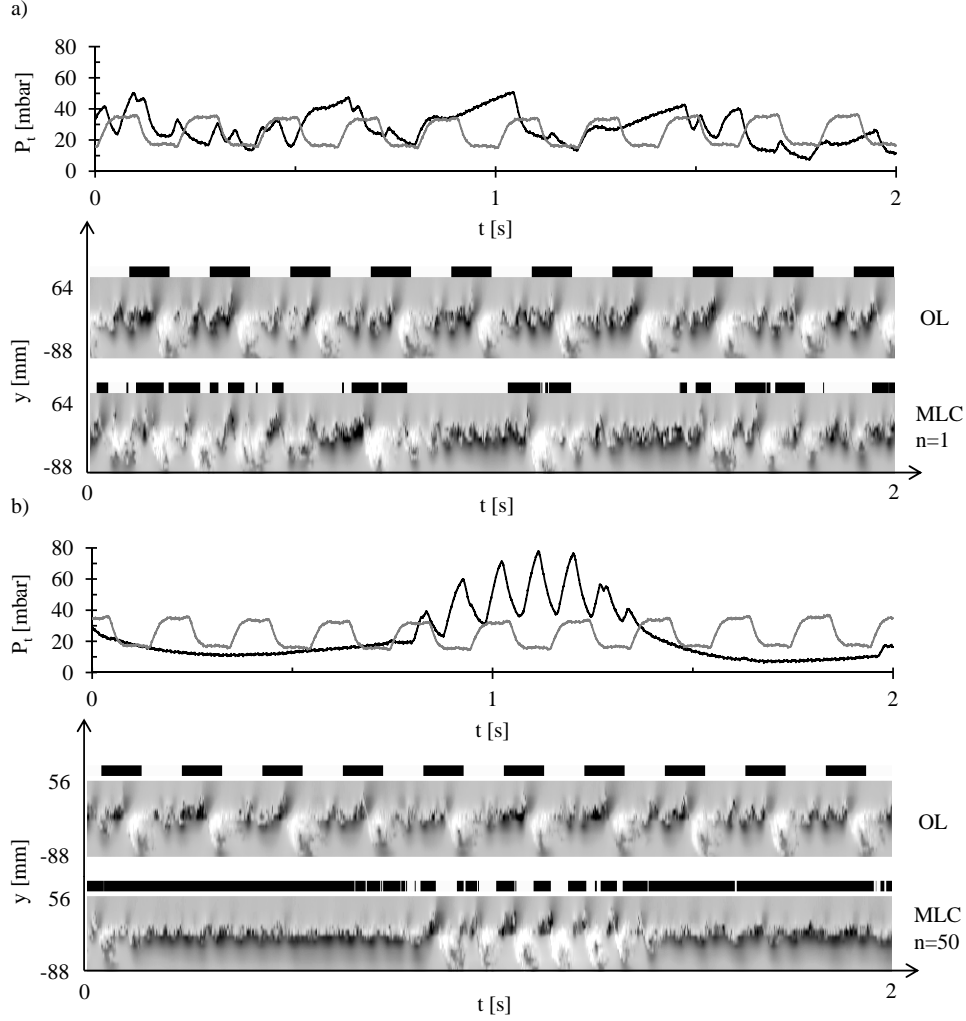


FIGURE 12. Results of MLC experiment with the high-gain amplitude controller for the open-loop reference and best MLC individual in generation a) $n = 1$ and b) $n = 50$. Top diagrams show instantaneous pressure in the plenum of the actuator system, for the open-loop case (grey line) and best MLC individual (black line). Bottom diagrams show corresponding pseudo-visualizations and actuation activity (black and white stripes). Greyscale map of visualization fields corresponds to u' in range $[-1.5, 1.5]$ (m s^{-1}).

”burst mode”), modulated by a very low frequency $f_{mod} < 1 \text{ Hz}$ (corresponding to the continuous blowing periods). The resulting cost function J_K values are $J_{K_{max}}^{OL} = 3.04$ and $J_{K_{max}}^{MLC} = 3.93$. In the case of this experiment a higher average amplitude of actuation was set for both control approaches ($P_t = 25 \text{ mbar}$), leading to momentum coefficient values of $C_{\mu}^{OL} = 0.58$ and $C_{\mu}^{MLC} = 0.89$. More information on the GP iterative process, leading to this solution, are given in section 5.1.

5. Discussion of MLC performance

In this section we discuss the convergence and capabilities of the MLC process in finding the best control law candidates. We explore the sensitivity of MLC to different parameters, such as: the size of the first and subsequent generations and settings for genetic operations. The statistically converged results are considered in the comparison between the performance of the best individuals and the reference open-loop actuations. Robustness to different flow conditions, for both open- and closed-loop control, is discussed. Finally, we examine the costs of MLC and other methods in terms of experimental time needed.

5.1. Influence of the variation of GP parameters

Two crucial questions naturally arise when dealing with any stochastic optimization method: i) what is the best obtainable solution and do we have any guarantee to obtain it, and ii) what are the optimal parameters for obtaining this solution? Due to the nature of the GP process, it is impossible to respond to the first question when a completely unknown system is concerned. In our case, however, we have at least a partial response in the form of the reference open-loop actuation. Parameters of GP we selected to vary were the population size and the balance of probabilities assigned to each of the three genetic operations.

The various effects of different GP parameters are sampled in figure 13. Each of the diagrams in this figure shows values of cost function J_K or J_W for open-loop reference (white circles) and best individual (black circles) for each generation. Best MLC individual data points are averaged values of the objective function for each instance this particular individual has been tested. This means that if an individual appeared in several generations (regardless of whether it was the best in every generation) all of its evaluations contribute to the value, similar to what is used to calculate the cost function in MLC as presented in section 3.3. The associated error bars denote the standard dispersion (assuming a Gaussian distribution) equal to three times the root-mean-square (*rms*) value of the best individual. Open-loop data points have no error bars since they are only evaluated once per generation. The dispersion of this data can be directly inferred from the diagrams.

An MLC experiment using the standard parameter set (see table 1) with a large initial and smaller subsequent population sizes is depicted in figure 13(a). It is apparent that the initial generation yields already the best individual, after which the GP process is only confirming the solution. An alternative experiment is shown in figure 13(b) where a sub-optimal solution is found in the first generation and it never evolves to a better one. This experiment used a small population size of 100 individuals in each generation, including the first. These results suggest that a Monte-Carlo random process (Metropolis & Ulam 1949), in forming the first generation, can be used to obtain the best solution if the population is large enough. An experiment involving 500 individuals per generation is shown in figure 13(c). Here, the best solution is also found in the first generation. This indicates that a population size of $N^{(1)} = 500$ individuals is enough to yield a solution by random probing of the search space, in this experiment.

Figure 13(d) shows the effect of a radical change in parameter settings. In this case, only 50 individuals were used in each generation, while genetic operations have been set to 0 % for replication, 65 % for mutation and 35 % for cross-over. Elitism is set to $N_e = 5$ best individuals. These settings assured that the initial generation had distinctly sub-optimal individuals, but gave ample room for evolution by introducing much more mutation, *i.e.* the ability to expand the search space beyond the possibilities of the first generation. A significant qualitative jump in cost function value of the best individual is

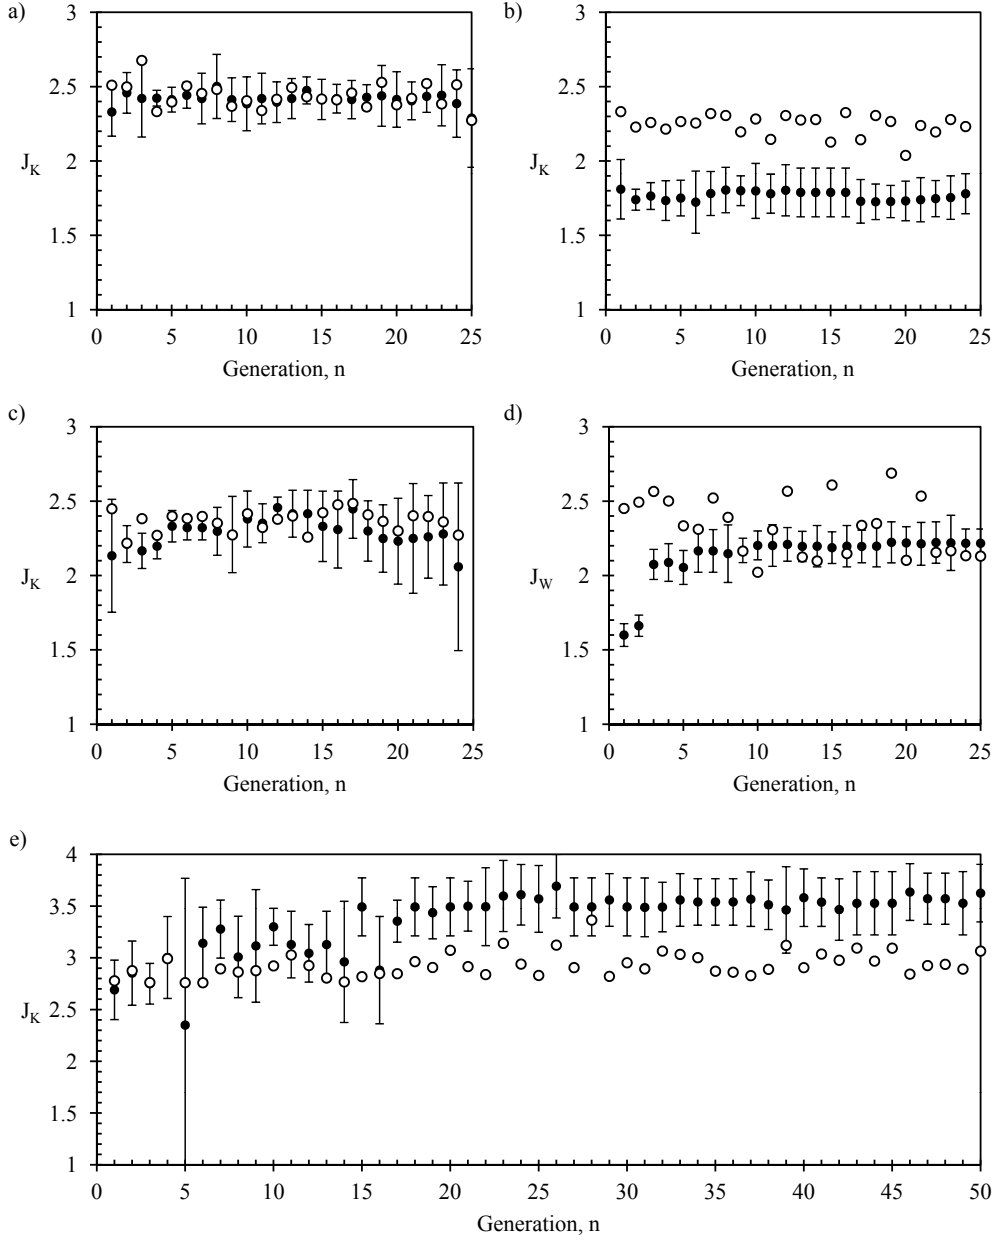


FIGURE 13. Convergence of MLC process using various GP and experimental parameters: a) MaxK experiment with standard GP settings (see table 1), b) MaxK experiment using a small population size ($N^{(n)} = 100$, for $n = 1, \dots, 25$), c) MaxK experiment using a medium population size ($N^{(n)} = 500$, for $n = 1, \dots, 25$), d) MaxW experiment using a very small population size and large mutation rate ($N^{(n)} = 50$, for $n = 1, \dots, 25$; $P_r = 0.0$, $P_x = 0.35$, $P_m = 0.65$), and e) MaxK experiment using the high-gain actuation amplitude controller ($N^{(n)} = 500$, for $n = 1, \dots, 50$).

apparent as early as the 3rd generation. This kind of behavior is a signature of a successful GP process. Testing time for this experiment until a stable best solution (which appears in generation $n = 6$) is significantly shorter than any previously shown experiment.

Another clearly successful GP process is shown in figure 13(e), this time for the case of the high-gain amplitude controller. In this case, as explained in section 4.7, the complexity of the controlled plant was increased by the response of the actuation amplitude control system. In terms of GP parameters, this experimental run was not different from the run shown in figure 13(c), *i.e.* it used 500 individuals in each generation and a standard set of other genetic operation parameters. Likewise, we can see that the best solution in the first generation was approximately equal to best open-loop control, as expected when a large initial population size is used. However, the MLC process evolved a better solution (appearing in generation $n = 15$), which no single frequency periodic actuation could have provided.

5.2. Statistical analysis of MLC performance

When the data from the GP process presented in the previous section is analyzed, it can be seen that both open- and closed-loop actuation has a level of standard dispersion from which it is sometimes difficult to reach a clear conclusion of which one is better. This deviation in performance results from the natural perturbations in the wind tunnel and in the actuator system.

A more complete way of making a comparison between open- and closed-loop control in this experiment is to average out all the results across the experiment. Resulting mean and *rms* cost function values (\bar{J}_K , \bar{J}_W and $J_{K_{rms}}$, $J_{W_{rms}}$) are shown in table 3, for all experiments discussed in section 4. The cost function values for the open-loop reference are averaged across their single evaluations in every generation. The best MLC individual is chosen from the last generation, but the averaging is performed using each evaluation of this individual, from every generation in which it appeared previously. This procedure should be sufficient to ensure comparable statistical sample sizes.

The values in table 3 indicate the open-loop actuation as having around 5% of an average performance advantage, in most cases. Analysis of the *rms* values shows they are of the order of 5–10% of the mean value. Depending on the experiment, such a high standard deviation can appear in either open- or closed-loop control case. Most obvious examples being the open-loop control in the MaxW (LS) experiment, and MLC best individual in the MaxK (LS) experiment, both for $x = 200$ mm. Thus, a direct performance comparison between open- and closed-loop control is rendered inconclusive, considering relatively large standard deviations of the cost function values. It is clear, however, that both control types have a significant impact on the mixing layer.

5.3. Robustness of MLC performance

A series of experiments was performed to explore the sensitivity of open- and closed-loop control to different flow conditions. Best scenarios of both types of control for maximization of objective function W (see section 4.3) were tested in a simulated change of Re_θ of the mixing layer. The un-actuated flow is first evaluated in both LS ($Re_\theta = 500$) and HS ($Re_\theta = 2000$) conditions (shown in figure 14a), for the purpose of estimating the results as a normalized cost function J_W .

In the first test of robustness (shown in figure 14b and c), we consider LS as a starting condition. Optimal open-loop and MLC control, which were found for the LS configuration were tested, yielding $J_{W_{max}}^{OL}(LS) = 2.42$ and $J_{W_{max}}^{MLC}(LS) = 2.15$, respectively. Once these evaluations were complete, the wind-tunnel speed was changed to the

	$x = 200$ [mm]		$x = 500$ [mm]	
MaxW (LS)	\bar{J}_W	$J_{W_{rms}}$	\bar{J}_W	$J_{W_{rms}}$
Open-loop	2.32	0.20	1.62	0.09
MLC	2.20	0.08	1.55	0.18
MaxK (LS)	\bar{J}_K	$J_{K_{rms}}$	\bar{J}_K	$J_{K_{rms}}$
Open-loop	2.44	0.08	1.85	0.19
MLC	2.52	0.25	1.73	0.13
MinK (LS)	\bar{J}_K	$J_{K_{rms}}$	\bar{J}_K	$J_{K_{rms}}$
Open-loop	0.70	0.04	0.81	0.05
MLC	0.75	0.04	1.04	0.12
MaxW (HS)	\bar{J}_W	$J_{W_{rms}}$	\bar{J}_W	$J_{W_{rms}}$
Open-loop	1.86	0.05	-	-
MLC	1.75	0.08	-	-
MaxK (high-gain PI)	\bar{J}_K	$J_{K_{rms}}$	\bar{J}_K	$J_{K_{rms}}$
Open-loop	2.93	0.12	-	-
MLC	3.48	0.16	-	-

TABLE 3. Mean (\bar{J}) and standard deviation (J_{rms}) cost function values of the best open-loop control and MLC individuals, for all the experiments presented in section 4.

HS settings and the controls were re-evaluated. The resulting cost function values are: $J_{W_{max}}^{OL}(LS \rightarrow HS) = 1.32$ and $J_{W_{max}}^{MLC}(LS \rightarrow HS) = 1.58$.

The average total enhancement of the cost function, across the two flow conditions, can be computed as:

$$\langle J_{W_{max}} \rangle_{LS,HS} = \frac{1}{2} [J_{W_{max}}(LS) + J_{W_{max}}(LS \rightarrow HS)], \quad (5.1)$$

for both types of control. The average cost function yields $\langle J_{W_{max}}^{OL} \rangle_{LS,HS} = 1.87$ for open-loop and $\langle J_{W_{max}}^{MLC} \rangle_{LS,HS} = 1.865$ for MLC. We can conclude that when these two controls are engaged and a speed change occurs, both will yield a similar average enhancement of around 87% of objective function W , compared to the un-actuated flow.

The second robustness test (shown in figure 14d and e) is started with HS as the initial condition, followed by a test in the LS setting. Here we test open-loop and MLC controls which were found best in HS conditions. We note the cost function values: $J_{W_{max}}^{OL}(HS) = 1.91$ and $J_{W_{max}}^{MLC}(HS) = 1.81$ for the starting conditions, and $J_{W_{max}}^{OL}(HS \rightarrow LS) = 1.41$ and $J_{W_{max}}^{MLC}(HS \rightarrow LS) = 1.94$, for the changed conditions. The average cost function values in this case are: $\langle J_{W_{max}}^{OL} \rangle_{HS,LS} = 1.66$ for open-loop, and $\langle J_{W_{max}}^{MLC} \rangle_{HS,LS} = 1.875$ for MLC. In this reversed speed change test, MLC has kept its score of an average W enhancement of 87%, while open-loop yielded only 66%.

As seen in the previous section, open-loop performs slightly better than closed-loop when tested in the optimal conditions: $J_{W_{max}}^{OL}(LS)$ and $J_{W_{max}}^{OL}(HS)$ in figure 14(b) and (d), respectively. When the conditions change, however, closed-loop control retains much of its effectiveness. The pseudo-visualization in figure 14(c) shows that closed-loop (LS \rightarrow HS) adopts a global frequency of actuation very similar to the optimal open-loop in HS

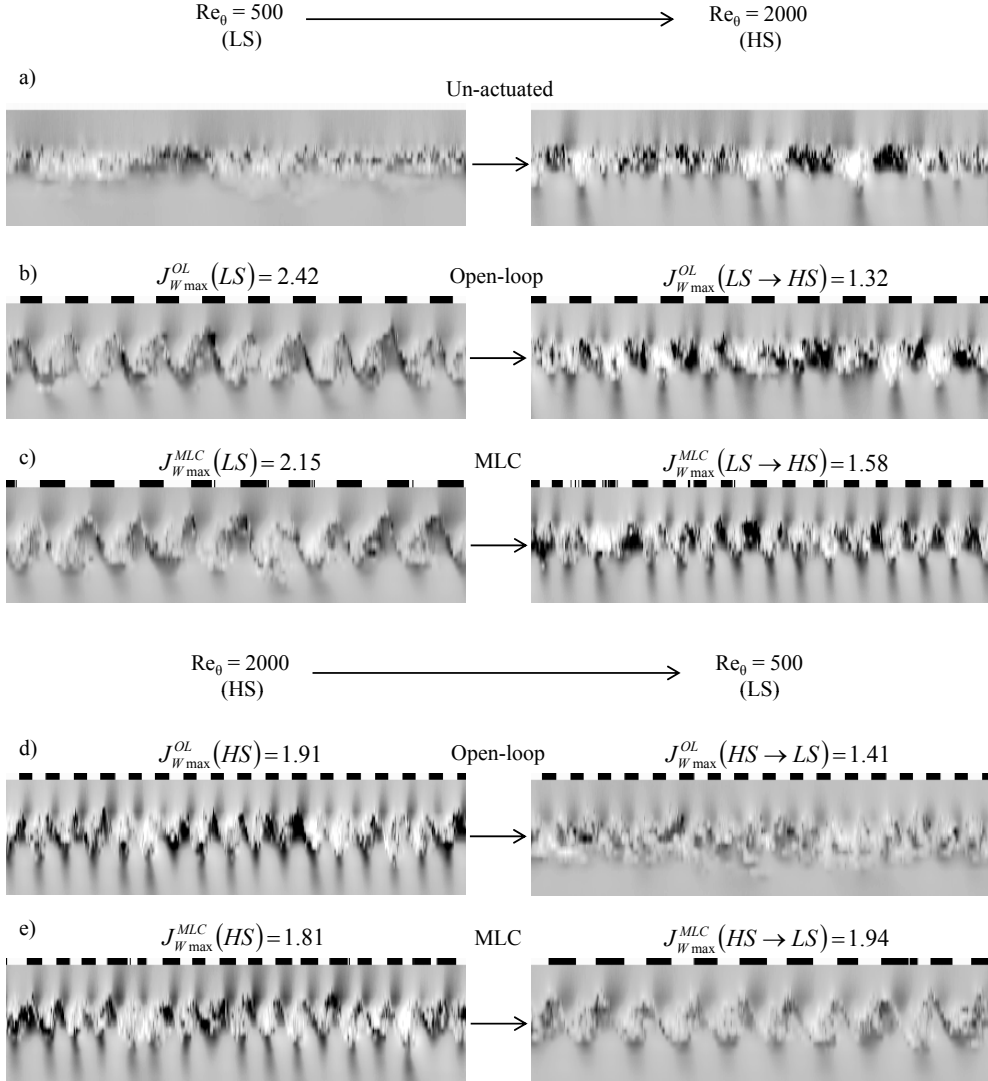


FIGURE 14. Robustness tests of open- and closed-loop control, with respect to two different speed configurations of the mixing layer. The best open-loop and the best MLC individual are tested in the conditions for which they are optimal, and then re-evaluated in a different configuration. Test of increasing Re_θ corresponds to: a) un-actuated, b) optimal LS open-loop, and c) best MLC for LS. Second test, decreasing Re_θ , is shown for: d) optimal open-loop for HS, and e) best MLC for HS. Pseudo-visualization is shown on a time scale of $t = 0.5$ s, but otherwise similarly as in figure 9.

configuration, shown in figure 14(d). A similar adaptation can be seen in figure 14(e), where the best MLC individual in HS \rightarrow LS configuration is very similar to the best open-loop actuation for LS conditions in figure 14(b).

We can also compare $J_{W_{\max}}^{OL}(LS \rightarrow HS)$ (figure 14b) to $J_{W_{\max}}^{OL}(HS)$ (figure 14d). Non-optimal open-loop scores 60% less W enhancement compared to what an optimal actuation would achieve. In a reverse test, the comparison of $J_{W_{\max}}^{OL}(HS \rightarrow LS)$ (figure 14d) and $J_{W_{\max}}^{OL}(LS)$ (figure 14b) yields a 100% reduction of effectiveness. MLC performs much




FIGURE 15. Estimation of experimental time needed for obtaining a best actuation solution using genetic programming, Monte-Carlo random process and open-loop mapping.

better; in similar comparisons for both tests, only 21–23% of reduction in effectiveness is recorded.

It can be concluded that open-loop can be a viable control method in varying flow conditions, only if it were to be applied using an adaptive control strategy (reference tracking or extremum seeking). Closed-loop control, on the other hand, shows an impressive ability to adapt and retain its effectiveness. Based on these results, a possible improvement of the MLC process could be to test multiple search spaces (based on flow conditions) in one experiment, in a search of an individual which scores the average best results across all conditions.

5.4. Time cost of MLC experiments

The study of the effects of parameter choice on the MLC convergence process reveals that very small population sizes are enough for the best solution to be obtained, and in just a few generations. These results have been repeated in both low- and high-speed mixing layer, and confirmed by several identical experimental runs.

The experiment shown in figure 13(d) indicates that the best solution can be seen as early as generation $n = 6$. As mentioned before, this experiment was performed using population sizes of $N = 50$ individuals in each generation. Consider that waiting for 2 more generations, until $n = 8$, is necessary to confirm if the solution is converged; this means that the total experimental time for obtaining this solution, taking into account 20 s for one evaluation, is $n \times N \times 20 \text{ s} = 2.2 \text{ h}$. This is substantially faster than evaluating a large first generation of 500 individuals and waiting for two more generations in order to confirm the best solution, which would be the minimum required to obtain the best solution through a Monte-Carlo process ($3 \times 500 \times 20 \text{ s} = 8.3 \text{ h}$).

Both of these experiments are, in turn, much shorter than a full mapping of the frequency/duty cycle space of periodic forcing available to the actuator system. Mapping would imply testing of frequencies up to 400 Hz in steps of 1 Hz and repeating the process using ten different duty cycles ($400 \times 10 \times 20 \text{ s} = 22 \text{ h}$). These projected time costs are comparatively shown in figure 15. If one adds testing in different flow configurations, the time costs would increase exponentially, putting the open-loop mapping at an even greater disadvantage. In this respect as well, MLC shows significant promise as a tool for exploring control of flow configurations whose dynamics are not well known in advance.

6. Conclusions and future directions

A model-free method for obtaining an optimal sensor-based control law for general multiple-input multiple-output (MIMO) plants, the “machine learning control” (MLC), is tested in an experimental mixing layer flow. MLC is based on genetic programming (GP), a function optimization method of machine learning. Randomly created linear or non-linear control laws are graded on-line with respect to a predefined cost function. Based on their success in minimizing or maximizing the cost function, the best are selected and evolved using genetic operations, to create the next generation of control law candidates. The process can be stopped when a satisfactory solution is found, or no further improvements are evident.

MLC has been successfully applied (Duriez *et al.* 2014) to a closed-loop stabilization of a generalized two-frequency mean-field model, and closed-loop control for the maximization of the Lyapunov exponent (stretching) of the forced Lorenz equations. In the former case, MLC allows detection and exploitation of frequency cross-talk, in an unsupervised manner. By definition, frequency cross-talk is ignored in any linearized system. In the case of the forced Lorenz equations, MLC provides an increase of unpredictability which is a highly nonlinear phenomenon. MLC is also applied in an experiment featuring the backward facing step flow (Gautier *et al.* 2014). Here, MLC yields a robust control law which causes a 50% reduction of the recirculation bubble, using a flapping mode manipulation mechanism.

In the mixing layer experiment, we first present the results of mapping of open-loop forcing as a reference for evaluation of the closed-loop control. The best MLC control laws behave similar to the best open-loop forcing, with respect to frequency and duty cycle properties of the actuation signal. Since these control laws are sensor-based, they cannot be as perfectly regular as periodic forcing. Hence, an open-loop actuation, optimally selected for a given sensor location and an objective function, has a better average performance. The periodic forcing of a mixing layer has been thoroughly studied in experiments (Oster & Wygnanski 1982; Fiedler 1998; Wiltse & Glezer 1993). These results all show that in order to change the local properties of a mixing layer, a clock-work type of actuation is sufficient. However, this clock-work needs to be precisely optimized for it to be effective at a given location in the convective flow of a mixing layer. This precise optimization is the weak point of open-loop actuation. By changing the initial Reynolds number of the mixing layer flow, the actuation parameters cease to be optimal and the impact of control is reduced.

Closed-loop control, on the other hand, can be expected to be robust to such changes. The best MLC control laws succeed not only in distilling the optimal clock-work actuation from the sensors in the mixing layer flow, but also in showing a remarkable ability to adapt to the changes of the initial conditions. The best control laws, found for one set of initial conditions, adopted the optimal actuation parameters when tested in very different conditions, for which they were not designed. MLC-based closed-loop control retains 80% of its effect for two extremes of the operational range of this wind tunnel. Compared to this, open-loop forcing is highly sensitive to such changes: its effect is reduced between 60% to 100% when conditions change.

MLC has a large chance to detect and exploit otherwise invisible local extrema. Such a case has been presented in section 4.7, where MLC successfully exploited the non-linear behavior of the actuation system to amplify the fluctuation levels in the mixing layer. This result is unobtainable by any periodic forcing, but it comes as a consequence of including (unwittingly) the amplitude controller as a variable of the controlled plant. In this case, MLC proves to be a good detector of the boundaries of the targeted plant.

The success of the MLC search/optimization process depends on the choice of parameters. Results show that a good direction to adopt is to start with a small population size and use aggressive settings for mutation. This corresponds to a fast, low-resolution sweep of a large search space, while still retaining the ability to optimize a good solution through cross-overs. This process is also proven to be far superior in experimental time needed, compared to a simple open-loop parametric study. Such an experiment is not very time-consuming and, therefore, several runs could be easily performed to confirm whether a solution is a stable one. A non-converging solution of 4-5 repetitions of this experiment, would indicate that larger population sizes should be used and different genetic programming settings. However, starting with larger population sizes might prove to be unnecessary, if the problem turns out to be simple enough. A random creation of the first generation individuals may already find the best solution, through a Monte-Carlo random process. Nevertheless, while a single generation can yield a best solution in such a way, one cannot know that this is the best solution. Results from several more generations would be needed, so that a convergence of the solution becomes obvious.

All the results regarding the MLC convergence process and the closed-loop control performance may be summed up as follows:

- Initial generation of random individuals can provide the best solution through a Monte-Carlo process, provided the problem at hand is simple enough and initial population size is large enough. In this case Genetic Programming can serve only to confirm the result or possibly optimize it.
- If the initial population does not yield the best solution, MLC, using a correct set of GP parameters, is able to achieve it in a very short time.
- For a given control problem, MLC will use any means provided to it to obtain a better solution. It is up to the user to provide a strong definition of the plant boundaries and minimize the possibility of external influences.
- Sensor-based control laws, using even the simplest forms of flow information from the mixing layer, can create an optimal actuation signal.
- Closed-loop control proves robust to changes in flow conditions.

We foresee that MLC can be significantly improved by adopting successful principles of control theory: Firstly, the sensor signals may be filtered so that they are less noise sensitive. Secondly, the actuation command may be included as argument in the control law. Thirdly, an ensemble of model-based control laws may constitute the first generation to be improved by MLC. Fourthly, the concept of reference tracking needs to be incorporated in future applications. Fifthly, time can become an argument of the control law, just like the sensor readings. In this case, MLC would be able to find the best open-loop forcing when closed-loop control is inevitably less effective. In addition, spatial combinations of the sensors, like in POD feedback control (Glauser *et al.* 2004; Parezanovic *et al.* 2014) may also provide a means to mitigate both noise and sensitivity to sensor placement.

The model-free formulation makes MLC a very flexible method. It can be applied to any MIMO plant and use any cost function formulation. Theoretically, no *a priori* knowledge of the plant is needed, but an MLC experiment using "fast" settings can be very useful for obtaining quick additional information about the flow. This information can then be used to improve the initial guess with respect to the underlying physics. Not only could this lead to better optimization of the control laws, it could also potentially reduce the total time needed to understand the dynamics of the plant.

Acknowledgements

The authors acknowledge the funding and excellent working conditions of the Senior Chair of Excellence 'Closed-loop control of turbulent shear flows using reduced-order models' (TUCOROM) supported by the French Agence Nationale de la Recherche (ANR) and hosted by Institute PPRIME. Additional funding by the ANR was provided via the SepaCoDe grant. Marc Segond would like to acknowledge the support of the LINC project (no. 289447) funded by EC's Marie-Curie ITN program (FP7-PEOPLE-2011-ITN). We thank the Ambrosys Ltd. Society for Complex Systems Management, the Bernd Noack Cybernetics Foundation and Hermine Freienstein-Witt for additional support.

We appreciate valuable stimulating discussions with the TUCOROM team (Jacques Borée, Eurika Kaiser, Nathan Kutz, Robert Niven, Michael Schlegel, Andreas Spohn and Gilles Tissot), with the SepaCoDe team, in particular Azeddine Kourta and Michel Stanislas, our PMMH collaborators Jean-Luc Aider and Nicolas Gautier, and Rudibert King, Christian Nayeri, Oliver Paschereit, Rolf Radespiel, Peter Scholz and Richard Semaan.

Special thanks are due to Nadia Maamar for a wonderful job in hosting the TUCOROM visitors.

REFERENCES

- BAGHERI, S., BRANDT, L. & HENNINGSON, D. S. 2009 Input-output analysis, model reduction and control of the flat-plate boundary layer. *J. Fluid Mech.* **620**, 263–298.
- BAGHERI, SHERVIN, MAZZINO, ANDREA & BOTTARO, ALESSANDRO 2012 Spontaneous symmetry breaking of a hinged flapping filament generates lift. *Physical review letters* **109** (15), 154502.
- BEAUDOIN, J. F., CADOT, O., AIDER, J. L. & WESFREID, J. E. 2006 Bluff-body drag reduction by extremum-seeking control. *Journal of Fluids and Structures* **22**, 973–978.
- BRUNTON, S. L., DAWSON, S. T. M. & ROWLEY, C. W. 2014 State-space model identification and feedback control of unsteady aerodynamic forces. *submitted for publication* p. <http://arxiv.org/pdf/1401.1473>.
- BRUNTON, S. L., ROWLEY, C. W. & WILLIAMS, D. R. 2013 Reduced-order unsteady aerodynamic models at low Reynolds numbers. *J. Fluid Mech.* **724**, 203–233.
- CATTAFESTA, III, L. N. & SHEPLAK, M. 2011 Actuators for Active Flow Control. *Annu. Rev. Fluid Mech.* **43**, 247–272.
- CHOI, H., MOIN, P. & KIM, J. 1994 Active turbulence control for drag reduction in wall-bounded flows. *J. Fluid Mech.* **262**, 75–110.
- CORKE, T. C., ENLOE, C. L. & WILKINSON, S. P. 2010 Dielectric Barrier Discharge Plasma Actuators for Flow Control. *Annu. Rev. Fluid Mech.* **42**, 505–529.
- DULLERUD, GEIR. E. & PAGANINI, FERNANDO 2000 *A course in robust control theory: A convex approach*. Berlin, Heidelberg: Springer.
- DURIEZ, T., PAREZANOVIĆ, V., CORDIER, L., NOACK, B. R., DELVILLE, J., BONNET, J.-P., SEGOND, M. & ABEL, M. 2014 Closed-loop turbulence control using machine learning. *Under consideration for publication in J. Fluid Mech.* .
- FIEDLER, H. E. 1998 Control of Free Turbulent Shear Flows. In *Flow Control: Fundamentals and Practices, Lecture Notes in Physics monographs*, vol. 53, pp. 335–429. Springer-Verlag, ISBN 3-540-63936-5.
- FLEMING, P. J. & PURSHOUSE, R. C. 2002 Evolutionary algorithms in control systems engineering: a survey. *Control engineering practice* **10** (11), 1223–1241.
- DE LA FRAGA, L. G. & TLELO-CUAUTLE, E. 2014 Optimizing the maximum Lyapunov exponent and phase space portraits in multi-scroll chaotic oscillators. *Nonlinear Dynamics* pp. 1–13.
- GAUTIER, N., AIDER, J.-L., DURIEZ, T., NOACK, B. R., SEGOND, M. & ABEL, M. 2014 Closed-loop separation control using machine learning. *Submitted to J. Fluid Mech. Available on arXiv.* .

- GLAUSER, M., HIGUCHI, H., AUSSEUR, J., PINIER, J. & CARLSON, H. 2004 Feedback control of separated flows. *AIAA paper* **2521**.
- GLEZER, A. & AMITAY, M. 2002 Synthetic jets. *Annu. Rev. Fluid Mech.* **34**, 503–529.
- GLEZER, A., AMITAY, M. & HONOHAN, A. M. 2005 Aspects of low-and high-frequency actuation for aerodynamic flow control. *AIAA journal* **43** (7), 1501–1511.
- GOLDBERG, D. 1989 *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Professional.
- HERVÉ, A., SIPP, D., SCHMID, P. J. & SAMUELIDES, M. 2012 A physics-based approach to flow control using system identification. *J. Fluid Mech.* **702**, 26–58.
- ILAK, MILOS 2009 Model reduction and feedback control of transitional channel flow. PhD thesis, Princeton University.
- ILAK, M. & ROWLEY, C. W. 2008 Modeling of transitional channel flow using balanced proper orthogonal decomposition. *Phys. Fluids* **20**, 034103.
- ILLINGWORTH, S. J., MORGANS, A. S. & ROWLEY, C. W. 2010 Feedback control of flow resonances using balanced reduced-order models. *Journal of Sound and Vibration* **330** (8), 1567–1581.
- JUANG, J. N. & PAPP, R. S. 1985 An eigensystem realization algorithm for modal parameter identification and model reduction. *Journal of Guidance, Control, and Dynamics* **8** (5), 620–627.
- JUANG, J. N., PHAN, M., HORTA, L. G. & LONGMAN, R. W. 1991 Identification of observer/Kalman filter Markov parameters: Theory and experiments. Technical Memorandum 104069. NASA.
- KERSTENS, W., PFEIFFER, J., WILLIAMS, D., KING, R. & COLONIUS, T. 2011 Closed-loop control of lift for longitudinal gust suppression at low Reynolds numbers. *AIAA Journal* **49** (8), 1721–1728.
- KING, R., ed. 2010 *Active Flow Control II, Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, vol. 108. Springer.
- KING, R., BECKER, R., FEUERBACH, G., HENNING, L., PETZ, R., NITSCHKE, W., LEMKE, O. & NEISE, W. 2006 Adaptive flow control using slope seeking. In *14th IEEE Mediterranean Conference on Control and Automation, 2006*, pp. 1–6.
- KOZA, J. R. 1992 *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. The MIT Press.
- KOZA, J. R., BENNETT III, F. H. & STIFFELMAN, O. 1999 *Genetic Programming as a Darwinian Invention Machine, Lecture Notes in Computer Science*, vol. 1598. Springer.
- LEWIS, M. A., FAGG, A. H. & SOLIDUM, A. 1992 Genetic programming approach to the construction of a neural network for control of a walking robot. In *IEEE International Conference on Robotics and Automation*, , vol. 3, pp. 2618–2623.
- LJUNG, L. 1999 *System Identification: Theory for the User*. Prentice Hall.
- LUCHTENBURG, D. M., GÜNTHER, B., NOACK, B. R., KING, R. & TADMOR, G. 2009 A generalized mean-field model of the natural and actuated flows around a high-lift configuration. *J. Fluid Mech.* **623**, 283–316.
- LUKE, S., PANAIT, L., BALAN, G., PAUS, S., SKOLICKI, Z., KICINGER, R., POPOVICI, E., SULLIVAN, K., HARRISON, J., BASSETT, J., HUBLEY, R., DESAI, A., CHIRCOP, A., COMPTON, J., HADDON, W., DONNELLY, S., JAMIL, B., ZELIBOR, J., KANGAS, E., ABIDI, F., MOOERS, H., O’BEIRNE, J., TALUKDER, K. A. & MCDERMOTT, J. 2013 A Java-based Evolutionary Computation Research System, <http://cs.gmu.edu/eclab/projects/ecj/>.
- MA, Z., AHUJA, S. & ROWLEY, C. W. 2011 Reduced order models for control of fluids using the eigensystem realization algorithm. *Theoretical and Computational Fluid Dynamics* **25** (1), 233–247.
- METROPOLIS, N. & ULAM, S. 1949 The monte carlo method. *Journal of the American Statistical Association* **44** (247), 335.
- MURPHY, K. P. 2012 *Machine Learning: A Probabilistic Perspective*. Cambridge: MIT Press.
- NORDIN, P. & BANZHAF, W. 1997 An On-Line Method to Evolve Behavior and to Control a Miniature Robot in Real Time with Genetic Programming. *Adaptive Behavior* **5** (2), 107–140.
- NORIEGA, J. R. & WANG, H. 1998 A direct adaptive neural-network control for unknown nonlinear systems and its application. *IEEE Trans. Neural Networks* **9** (1), 27–34.

- OSTER, D. & WYGNANSKI, I. 1982 The forced mixing layer between parallel streams. *J. Fluid Mech.* **123**, 91–130.
- PAEZANOVIC, V., LAURENTIE, J.-C., FOURMENT, C., DELVILLE, J., BONNET, J.-P., SPOHN, A., DURIEZ, T., CORDIER, L., NOACK, B. R., ABEL, M., SEGOND, M., SHAQARIN, T. & BRUNTON, S. L. 2014 Mixing layer manipulation experiment: From open-loop forcing to closed-loop machine learning control. *Submitted to Flow, Turbulence and Combustion*.
- PASTOOR, M., HENNING, L., NOACK, B. R., KING, R. & TADMOR, G. 2008 Feedback shear layer control for bluff body drag reduction. *J. Fluid Mech.* **608**, 161–196.
- ROWLEY, CLARENCE W. 2002 Modeling, simulation, and control of cavity flow oscillations. PhD thesis, California Institute of Technology.
- ROWLEY, C. W. 2005 Model reduction for fluids using balanced proper orthogonal decomposition. *International Journal of Bifurcation and Chaos* **15** (3), 997–1013.
- SAMIMY, M., DEBIASI, M., CARABALLO, E., SERRANI, A., YUAN, X., LITTLE, J. & MYATT, J. H. 2007 Feedback control of subsonic cavity flows using reduced-order models. *J. Fluid Mech.* **579**, 315–346.
- SEMERARO, O., BAGHERI, S., BRANDT, L. & HENNINGSON, D. S. 2013 Transition delay in a boundary layer flow using active control. *J. Fluid Mech.* **731**, 288–311.
- THIRIA, B., GOUJON-DURAND, S. & WESFREID, J.E. 2006 The wake of a cylinder performing rotary oscillations. *J. Fluid Mech.* **560**, 123–147.
- VUKASINOVIC, B., RUSAK, Z. & GLEZER, A. 2010 Dissipative small-scale actuation of a turbulent shear layer. *J. Fluid Mech.* **656**, 51–81.
- WAHDE, M. 2008 *Biologically Inspired Optimization Methods: An Introduction*. WIT Press.
- WILTSE, J. M. & GLEZER, A. 1993 Manipulation of free shear flows using piezoelectric actuators. *J. Fluid Mech.* **249**, 261–285.

Appendix A. Linear System Identification and Models

Here, we investigate linear models based on input–output data from the mixing layer experiment. First, we introduce the particular methods for system identification employed for this data. We then present results demonstrating the performance of linear models for predicting relevant quantities, such as a local turbulent energy content. It is demonstrated that linear models are unable to capture significant flow features of the strongly nonlinear mixing layer.

A.1. ERA/OKID approach

To provide context for the unsupervised machine learning methods in this work, we first investigate the performance of linear models of the strongly nonlinear mixing layer. In particular, we construct linear input–output models from impulse response data using the eigensystem realization algorithm (ERA) of Juang & Pappa (1985). ERA models have recently been shown to be equivalent to those obtained using balanced proper orthogonal decomposition (BPOD) (Rowley 2005; Ma *et al.* 2011), but without the need for adjoint simulations. Therefore, ERA models of a given reduced order r will capture the r most observable and controllable fluid states in a balanced representation. Balanced models are particularly useful for control since they faithfully capture the input–output relationship for a particular flow configuration and set of sensors and actuators (Rowley 2005; Ilak & Rowley 2008; Bagheri *et al.* 2009; Illingworth *et al.* 2010). Upper and lower error bounds on model performance are available in terms of the Hankel singular values when ERA/BPOD is used to reduce the model order of a linear system (Rowley 2005; Dullerud & Paganini 2000); this is due to a connection between ERA, BPOD, and balanced truncation. Although we use ERA, there are many alternative methods to identify state-space models from input–output data (Ljung 1999).

ERA is considered a *system identification* technique, since it is based on measurement data and does not rely on knowledge of an underlying model. Given output measure-

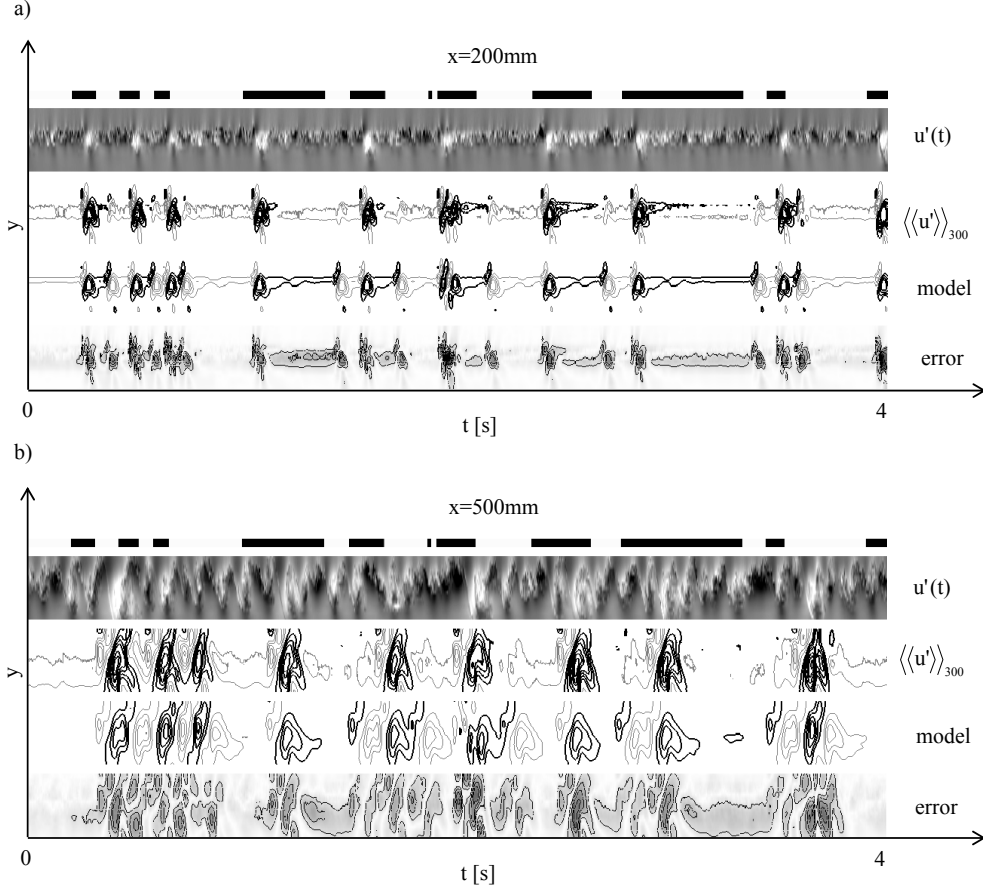


FIGURE 16. Pseudo-visualization of the first four seconds of (top-down): the actuation signal, $u'(t)$, $\langle u' \rangle_{300}$, model reconstruction and error, for a) $x = 200 \text{ mm}$ and b) $x = 500 \text{ mm}$.

ments for an impulse response in the actuation, ERA will produce a balanced state-space model of a desired model order r . However, in many fluid experiments the impulse response data is noisy, and the underlying system may be strongly nonlinear. In this case, we use frequency-rich input actuation sequences in conjunction with the observer/Kalman filter identification (OKID) (Juang *et al.* 1991), which provides a cleaner estimate of the impulse response despite noisy measurements (Brunton *et al.* 2013, 2014). This filtered impulse response is then passed through ERA to obtain a reduced-order model.

A.2. Experimental model ID

For the mixing layer experiment, we implement a pseudo-random actuation sequence where the jets are turned on and off for random intervals τ over a 30 second interval. The variable τ is sampled from a Poisson distribution with $\lambda = 4$ and a mean hold-time of $\tau = 0.1$ seconds. Hot-wire measurements (see section 2) are collected for 300 repeated experimental runs and phase averaged, with respect to the actuation signal, to reduce irregular transient events. The idea of using a pseudo-random pulsed-blowing signal is inspired by Kerstens *et al.* (2011).

Before passing the phase-averaged data through OKID, it is first coarsened temporally

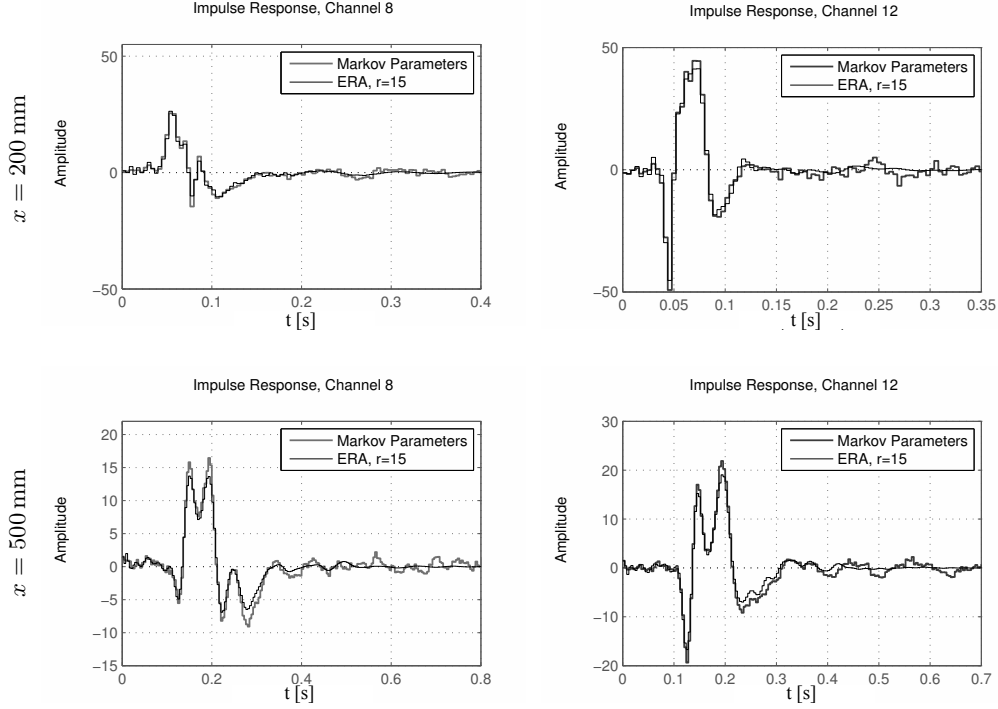


FIGURE 17. Markov parameters from OKID and impulse response of ERA model.

to $\Delta t = 0.004 s$. Next, since the downstream mixing layer is in a steady-state of saturated nonlinearity, corresponding to broad-band oscillations, the OKID filtered impulse response is multiplied by a hyperbolic tangent function that starts at 1 and decays to 0; this is critical to ensure a linearized impulse response that decays, resulting in stable models. We then generate ERA models of order $r = 15$ for each of the phase-averaged hot-wire signals; $r = 15$ provides a balance between accuracy and model order, and results are qualitatively similar for $r = 10$ and $r = 20$.

Figure 16 shows the actuation signal, an instance of the fluctuations $u'(t)$, the phase-averaged signal $\langle\langle u' \rangle\rangle_{300}$, ERA model reconstruction and error for locations at $x = 200 mm$ and $x = 500 mm$. From this figure, it is clear that the downstream velocity measurements are strongly correlated with the actuation signal, and that velocity fluctuation is enhanced. Comparing the phase-averaged measurements $\langle\langle u' \rangle\rangle_{300}$ to a single instance u' , it is apparent that phase-averaging filters out many important transient flow nonlinearities, resulting in significantly decreased fluctuation levels. Therefore, even a perfect model reconstruction of $\langle\langle u' \rangle\rangle_{300}$ would only capture a small portion of the true instantaneous fluctuation energy. Moreover, the model reconstruction, although qualitatively similar to the phase-averaged measurements, has significant error on the same order of magnitude as the phase-averaged measurements.

Figure 17(top) shows the Markov parameters (output of OKID) and the ERA reconstruction for hot-wire channels 8 and 12 at $x = 200 mm$. Although the Markov parameters are well approximated by the ERA model, there is still significant reconstruction error when these models are used to approximate the original phase-averaged data, as shown in figure 16. This suggests that even the phase-averaged measurements exhibits strongly nonlinear responses to actuation. Indeed, there are two important types of nonlinearity for this system. The first type of transient, irregular flow events are removed




FIGURE 18. Bode plot for all hot-wire signals at a) $x = 200\text{ mm}$ and b) $x = 500\text{ mm}$. Hot-wire channels 1–10 are indicated by solid lines, and 11–19 are dashed. The lines are darker towards the middle of the rake and lighter toward the edges.

by the phase-averaging procedure. The second type of nonlinearity consists of regular, repeatable responses to actuation, which persist in the phase-averaged measurements. The linear ERA models are incapable of capturing these nonlinearities, and therefore only capture a small portion of the phase-averaged response. These two losses greatly diminish the effectiveness of linear models for predicting the nonlinear flow in response to actuation. Similar plots of the Markov parameters, reconstruction, and error are shown for $x = 500\text{ mm}$ in figure 17(bottom) and figure 16.

The frequency responses (i.e., Bode plots) of the ERA models for all hot-wire channels are shown in figure 18 for the $x = 200\text{ mm}$ and $x = 500\text{ mm}$ locations. First, notice the presence of broadband resonant behavior in the middle-frequency range, as well as the presence of a small feed-through term, indicated by the high-frequency asymptote of the magnitude plot. As the hot-wire rake is moved downstream, the feed-through term reduces significantly, as expected. The resonant phenomenon also becomes more coordinated, and the phase plot indicates a larger time-delay. From the magnitude plots, it is possible to estimate the energy-based cost function J_K at various frequencies by summing the square of the magnitude across all hot-wires. The plots of this reconstruction of the cost function J_{model} are shown in figure 19 for the $x = 200\text{ mm}$ and $x = 500\text{ mm}$ locations. These predictions correspond to the case of 50% duty cycle actuation since both the “on” and “off” periods in the input actuation signal were sampled from the same Poisson distribution. Comparing with the experimental values from figure 8 at low frequencies, the linear model prediction of the cost function J_{model} has similar features, but the amplitude of fluctuation energy is off by a factor of three. Moreover, the linear models incorrectly predict zero velocity fluctuation energy at high frequencies.

There are a number of important conclusions that may be drawn from this linear

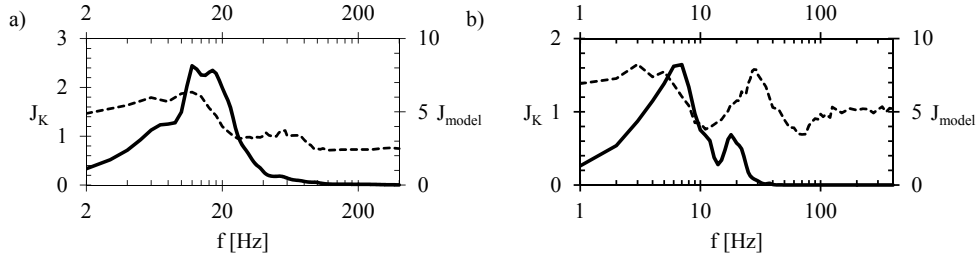


FIGURE 19. Prediction of the mixing layer response to open-loop actuation according to the linear model (solid line), compared to the experiment (dashed line), for a) $x = 200 \text{ mm}$ and b) $x = 500 \text{ mm}$.

model identification. First, it is clear that linear models obtained from the ERA/OKID procedure on phase-averaged hot-wire measurements have significant errors, on the same order as the averaged measurements. The close agreement between Markov parameters and the ERA model indicates that this error is due to nonlinear flow responses that persist in the phase-averaged signal. Moreover, the phase-averaged measurements fail to capture a significant portion of the turbulent fluctuations seen in a single instance of the response to actuation. These compound errors highlight the severe limitations of linear systems to characterize the strongly nonlinear mixing layer. However, the frequency response and predicted J_{model} do capture some important qualitative features.